Dynamic Asset Allocation with Transaction Costs:

The Importance of Hedging Demands

Pierre Collin-Dufresne Kent Daniel Mehmet Saglam

September 2019

Large Literature on 'Frictionless' Dynamic Portfolio Choice

- Markowitz's (1958) one-period mean-variance efficient (MVE) portfolio choice is still widely popular among practitioners.
- Merton (1969,1971) and Cox-Huang (1989) introduce **Hedging Demand**:

→ In dynamic setting, it may be optimal to deviate from instantaneous MVE portfolio to hedge against future changes of intercept (risk-free rate) and slope (sharpe ratio) of the conditional MVE frontier.

- Lots of academic work on **strategic asset allocation** show empirical relevance of hedging demand:

→ Return predictability significantly affects long-term asset allocation. (Brennan, Schwartz, Lagnado (1997), Brandt (1999), Kim and Omberg (1996), Campbell (1999) and Campbell and Viceira (2002)...)

Q? Insights from that literature seems to be mostly ignored by practitioners (deliberately?).

→ Realistic asset allocation needs to impose realistic objective function, constraints, transaction and price impact ('slippage') costs.

Traditional Literature on Dynamic Portfolio Choice with T-costs

- Constantinides (1986) Davis-Norman (1990) Dumas-Luciano (1991) derive optimal (No-trade region) policy under proportional costs in i.i.d. return setting.
  - Requires sophisticated mathematics (e.g., 'viscosity solutions' Shreve-Soner (1994)) even though limited to one risky asset with i.i.d. returns.
  - Multi-asset analysis in i.i.d. case limited to CARA preferences (Liu (2004)).
- Studies combining predictability in returns and t-costs typically limited to one or two risky assets and use numerical solutions (Balduzzi and Lynch (1999), Lynch and Tan (2011), Longstaff (2001)).
- Optimal execution of fixed size position to trade-off transaction costs and risk (Almgren and Chriss (2000), Grinold and Kahn (chap 16))

The Linear Quadratic Framework

- In a Linear-Quadratic (LQ) framework explicit solution for optimal portfolio choice for many stocks, with many predictors, and quadratic t-costs (i.e., linear price impact) obtains
  (Litterman (2005), Grinold (2006), Garleanu-Pedersen (2013), Collin-Dufresne, Daniel, Saglam (2018))

- The LQ framework relies on ad-hoc instantaneous mean-variance objective function which is myopic in the absence of t-costs.

- This paper:
  - → Standard preferences which nest (micro-found) this LQ-objective function.
  - → Explicit portfolio choice solution for (i) general (non-myopic) objective (ii) arbitrary number or assets, (iii) predictable returns.
  - → How important are hedging demands for portfolio choice with transaction costs?

Asset allocation with Transaction costs and predictability

- What do (quant) practitioners do?
  - Identify return forecasting factors
  - Identify risk factors
  - Create portfolios with 'optimal' exposures to these factors that
  - Maximize expected returns, net of expected trading costs, subject to a risk budget
- As an example consider optimal combination of three signals for large cross-section of individual stocks
  - Short term reversal (REV: half-life of 5 days)
  - Momentum (MOM: half-life 150 days)
  - Value (VAL: long-term reversal with a half-life of 700 days)
- Each stock will have a specific exposure to REV, MOM, and VAL. These factors will decay at different rates and are clearly not independent.
- How do we operationalize the 'optimal trade-off' when there are transaction costs?
  - If trade more often, expect to capture more alpha, but pay more transaction costs.
  - If trade less often, may not benefit from fast signals.

  → Should we trade more/less aggressively when signals decay faster/slower?
  - Trading more frequently reduces tracking error, but increases t-costs.

  → Should we trade more/less aggressively when volatility is larger/smaller?

A one-period (practitioner) approach

- A one period problem with quadratic transaction costs:

$$\max_{n_{t+1}} \left\{ n_{t+1}^\top \mu - h \underbrace{\frac{1}{2}(n_{t+1} - n_t)^\top \Lambda (n_{t+1} - n_t)}_{t-costs} - \frac{\gamma}{2} n_{t+1}^\top \Sigma\, n_{t+1} \right\}$$

- where
  - n_t is vector of **initial** stock holdings
  - μ is vector of **expected** returns
  - Σ is covariance matrix of returns
  - Λ is price impact matrix ($\sim$ quadratic transaction costs).
- If $\Lambda = \lambda\Sigma$ the optimal position is

$$n_{t+1}^* = (1-\tau)\, n_t + \tau\, (\gamma\Sigma)^{-1} \mu$$

$\Rightarrow$ Trade-off t-costs vs. tracking error by partially trading towards Markowitz portfolio with trading speed:

$$\tau = \frac{1}{1 + h\lambda/\gamma}$$

- Practitioners typically:
  - target a level of risk with risk-aversion γ,
  - calibrate to match realistic price impact Λ estimates, and
  - optimize over the trading speed parameter (h) to maximize backtest performance.

The dynamic model setup and Linear-Quadratic framework

- N-vector of stock price S_t has dynamics:

$$dS_t = (\mu_0 + \mu x_t)dt + \sigma_s dZ_s(t) \quad (1)$$
$$dx_t = -\kappa x_t dt + \sigma_x dZ_x(t) + \sigma_{xs} dZ_s(t) \quad (2)$$

N.B.: nests the case where every stock expected return is driven by K stock-specific predictors with different decay rates (e.g., momentum and value).

- Assuming risk-free $r = 0$, wealth dynamics with position vector n_t:

$$dW_t = n_t^\top dS_t - \frac{1}{2}\theta_t^\top \Lambda \theta_t dt \quad (3)$$
$$= n_t^\top (\mu_0 + \mu x_t)dt + n_t^\top \sigma_s dZ_s(t) - \frac{1}{2}\theta_t^\top \Lambda \theta_t dt \quad (4)$$
$$dn_t = \theta_t dt \quad (5)$$

- Classic but *ad-hoc* instantaneous mean-variance (LQ) objective function:

$$\max_\theta \mathrm{E}\left[\int_0^T \left\{dW_t - \frac{1}{2}\gamma dW_t^2\right\}\right] = \max_\theta \mathrm{E}\left[\int_0^T \left\{n_t^\top(\mu_0 + \mu x_t) - \frac{1}{2}\theta_t^\top \Lambda \theta_t - \frac{1}{2}\gamma n_t^\top \Sigma n_t\right\} dt\right]$$

where $\Sigma = \sigma_s \sigma_s^\top$ is the covariance of returns.
(widely used: Litterman (2005), Grinold (2006), Garleanu-Pedersen (2013), Collin-Dufresne, Daniel, Saglam (2019), Duffie-Zhu (2017), Vives (2011)...)

The general objective function: finite horizon case

- Consider solution $(H_t, \sigma_{H,s}, \sigma_{H,x})$ to the recursive equation:

$$H_t = \mathrm{E}_t \left[W_T - \int_t^T \left\{ \frac{1}{2}\gamma ||\sigma_{H,s}||^2 + \frac{1}{2}\gamma_x ||\sigma_{H,x}||^2 \right\} du \right] \tag{6}$$

- H_t is certainty equivalent of source-dependent stochastic differential utility agent, with absolute risk-aversion coefficient γ towards Z_s and γ_x towards Z_x shocks.
  - When $\gamma_x = \gamma$ it nests CARA expected utility:

$$H_t = -\frac{1}{\gamma} \log(\mathrm{E}_t[e^{-\gamma W_T}]). \tag{7}$$

  - When $\gamma_x \sigma_x \to 0$ and $\sigma_{xs} = 0$ it nests instantaneous mean-variance:

$$H_t = W_t + \mathrm{E}_t \left[\int_t^T \left\{ dW_u - \frac{1}{2}\gamma dW_u^2 \right\} \right]. \tag{8}$$

$\rightarrow$ The instantaneous mean-variance investor is risk-neutral with respect to changes in the investment opportunity set. She only displays risk-aversion towards 'level' shocks.

Q? Compare solution to (7) with solution to (8) to understand impact of non-myopic (e.g., hedging demand) on portfolio choice with transaction costs.

The Optimal Portfolio rule

- If $\Lambda = 0$ (no t-costs) the optimal position (Merton 1971):

$$n_t = \underbrace{(\gamma\Sigma)^{-1}(\mu_0 + \mu x_t)}_{Markowitz} - \underbrace{\Sigma^{-1}\Sigma_{sx}(c_1(t) + c_2(t)x)}_{HedgingDemand} \tag{9}$$

where $\Sigma = \sigma_s \sigma_s^\top$ and $\Sigma_{sx} = \sigma_s \sigma_{xs}^\top$.

$\Rightarrow$ If zero correlation between returns and expected returns (i.e., $\Sigma_{sx} = 0$) it is optimal to hold the instantaneous mean-variance efficient Markowitz portfolio.

- If $\Lambda > 0$ positive definite, then

$$dn_t = \tau_t(aim(x_t, t) - n_t)\, dt \tag{10}$$

$$\tau_t = \Lambda^{-1} Q(t) \tag{11}$$

$$aim(x, t) = Q(t)^{-1}(q_0(t) + q(t)^\top x) \tag{12}$$

where Q, q, q_0 solve system of Riccatti-style ODE and $aim(t, x) = \operatorname{argmax}_n J(n, x, t)$.

Insights from the solution: the instantaneous mean-variance case

- For an instantaneous mean-variance agent (eq. (8)), the optimal trading speed depends only on the eigenvalues (η_i) and eigenvectors (F_i) of $\gamma\Lambda^{-1}\Sigma \equiv FD_\eta F^\top$:

$$\tau_t = FD_h(t)F^\top \quad \text{with} \quad h_i(t) = \sqrt{\eta_i}\frac{1 - e^{-2\sqrt{\eta_i}(T-t)}}{1 + e^{-2\sqrt{\eta_i}(T-t)}}$$

- The aim portfolio can be written as $aim(x,t) = (\gamma\Sigma)^{-1}\int_t^T \omega_{t,u}\mu_S(t,u)du$, where:

$$\mu_S(t,u) = \frac{1}{dt}\mathrm{E}_t[dS_u] = \mu_0 + \mu e^{-\int_t^u \kappa ds}x_t$$

$$\omega_{t,u} = (\int_t^T e^{-\int_t^z \tau_s^\top ds}dz)^{-1}e^{-\int_t^u \tau_s^\top ds}$$

- As in one-period mean-variance benchmark, it is optimal to "trade partially towards aim portfolio", but:
  - Trading speed is endogenous and depends only on $\gamma\Lambda^{-1}\Sigma$.
  - Aim portfolio equals the Markowitz MVE portfolio if $\Lambda = 0$ **or** $\kappa = 0$.
  - Else, Aim portfolio $\sim$ *Markowitz* which replaces mean with expected future returns discounted for (i) trading speed (τ), and (ii) signal mean-reversion (κ).
  - Allocation is independent of the correlation between returns and signals (Σ_{sx}).

Results extend to Random Horizon (stationary solution)

- Consider solution $(H_t, \sigma_{H,s}, \sigma_{H,x})$ to the recursive equation with random horizon $\mathcal{T}$ (Poisson with intensity ρ):

$$H_t = \mathrm{E}_t\left[W_{\mathcal{T}} - \int_t^{\mathcal{T}} \left\{\frac{1}{2}\gamma||\sigma_{H,s}||^2 + \frac{1}{2}\gamma_x||\sigma_{H,x}||^2 + \rho\left(W_s - H_{s-} - \frac{1-e^{-\gamma_{\mathcal{T}}(W_s - H_{s-})}}{\gamma_{\mathcal{T}}}\right)\right\} ds\right]$$

- Then H_t is the certainty equivalent of source-dependent SDU agent with ARA coefficient γ toward Z_s, γ_x towards Z_x, and $\gamma_{\mathcal{T}}$ towards horizon arrival $\mathbf{1}_{\{\mathcal{T}\leq t\}}$.
  - When $\gamma_{\mathcal{T}} = \gamma_x = \gamma$, it nests CARA expected utility:

$$H_t = -\frac{1}{\gamma}\log(\mathrm{E}_t[e^{-\gamma W_{\mathcal{T}}}]). \tag{13}$$

  - When $\gamma_{\mathcal{T}} = \gamma_x\sigma_x \to 0$ and $\sigma_{xs} = 0$, it nests the discounted instantaneous mean-variance objective function:

$$H_t = W_t + \mathrm{E}_t\left[\int_t^{\infty} e^{-\rho(u-t)}\left\{dW_u - \frac{1}{2}\gamma dW_u^2\right\}\right]. \tag{14}$$

$\rightarrow$ Instantaneous mean-variance investor is risk-neutral with respect to both shocks to expected returns and horizon risk.

- In the following we focus on stationary solution case with $\gamma_{\mathcal{T}} = 0$ (i.e. we ignore horizon risk premium) and compare solution with $\gamma_x = 0$ to with $\gamma_x = \gamma > 0$ to investigate role of the hedging demand.

The Optimal Portfolio rule

- If $\Lambda = 0$ (no t-costs) the optimal position (Merton):

$$n_t = \underbrace{(\gamma\Sigma)^{-1}(\mu_0 + \mu x_t)}_{Markowitz} - \underbrace{\Sigma^{-1}\Sigma_{sx}(c_1 + c_2 x)}_{HedgingDemand} \tag{15}$$

$\Rightarrow$ If zero correlation ($\Sigma_{sx} = 0$) optimal to hold the instantaneous mean-variance markowitz portfolio:

- If $\Lambda > 0$ positive definite, then

$$dn_t = \tau_t(aim(x_t) - n_t)\, dt \tag{16}$$

$$\tau_t = \Lambda^{-1} Q \tag{17}$$

$$aim(x, t) = Q^{-1}(q_0 + q^\top x) \tag{18}$$

where Q, q, q_0 solve system of equation and $aim(x) = \mathrm{argmax}_n\, J(n, x)$.

Insights from the solution: the instantaneous mean-variance case

- For an instantaneous mean-variance agent, the optimal trading speed matrix $\tau = \Lambda^{-1}Q$ is given by:

$$\tau = FD_hF^\top$$
$$h_i = \frac{1}{2}(\sqrt{\rho^2 + 4\eta_i} - \rho)$$

where (F_i, η_i) are eigenfactors and eigenvalues of $\gamma\Lambda^{-1}\Sigma$.

- The optimal aim portfolio

$$aim(x_t) = Q^{-1}(q_0 + q^\top x_t) \tag{19}$$
$$= (\gamma\Sigma)^{-1}\int_0^\infty \omega_u \mathrm{E}_t[\mu_S(x_{t+u})]du \tag{20}$$
$$\mu_S(x_t) = \frac{1}{dt}\mathrm{E}_t[dS_t] = \mu_0 + \mu x_t \tag{21}$$
$$\omega_u = (\rho + \tau^\top)e^{-(\rho+\tau^\top)u} \tag{22}$$

→ Similar to finite horizon solution, but stationary

Illustration of the trading rule in the one-stock & one factor case

When trading costs are small: Merton hedging demand depends on $\sigma_{xs} \lesseqgtr 0$

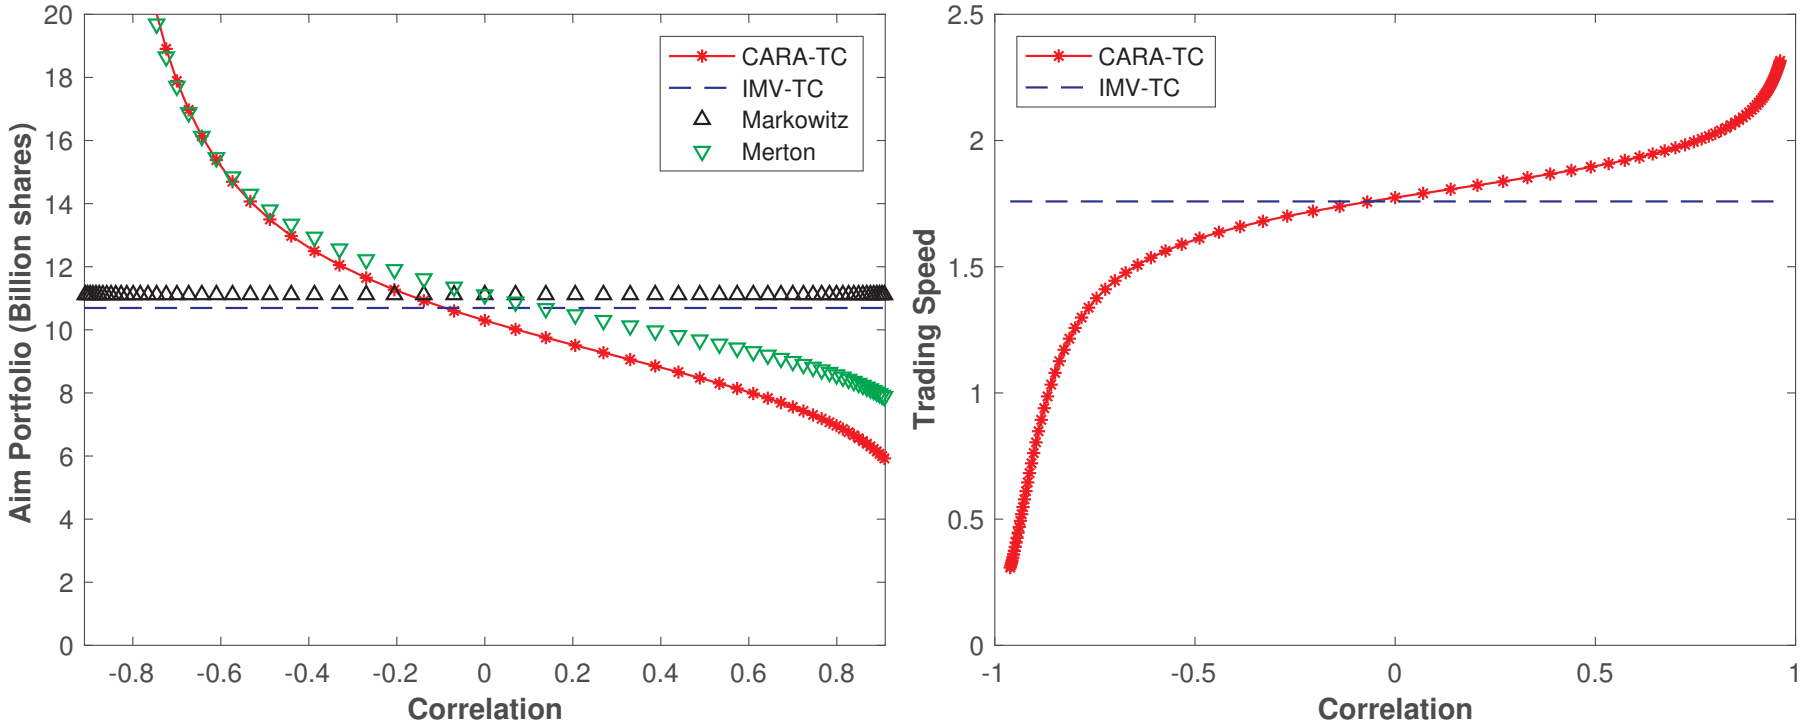

Figure: Parameters: $\mu_0 = 0$, $\mu = 1$, $\kappa = 0.1$, $\sigma_s = 0.3$, $\sigma_x = 0.1$, $\Lambda = 2 \times 10^{-11}$, $\gamma_x = 10^{-9}$, $\gamma = 10^{-9}$, $x_0 = 1$, $\rho = 0.8$.

For higher t-costs, hedging demands are very important (especially for $\sigma_{xs} < 0$)

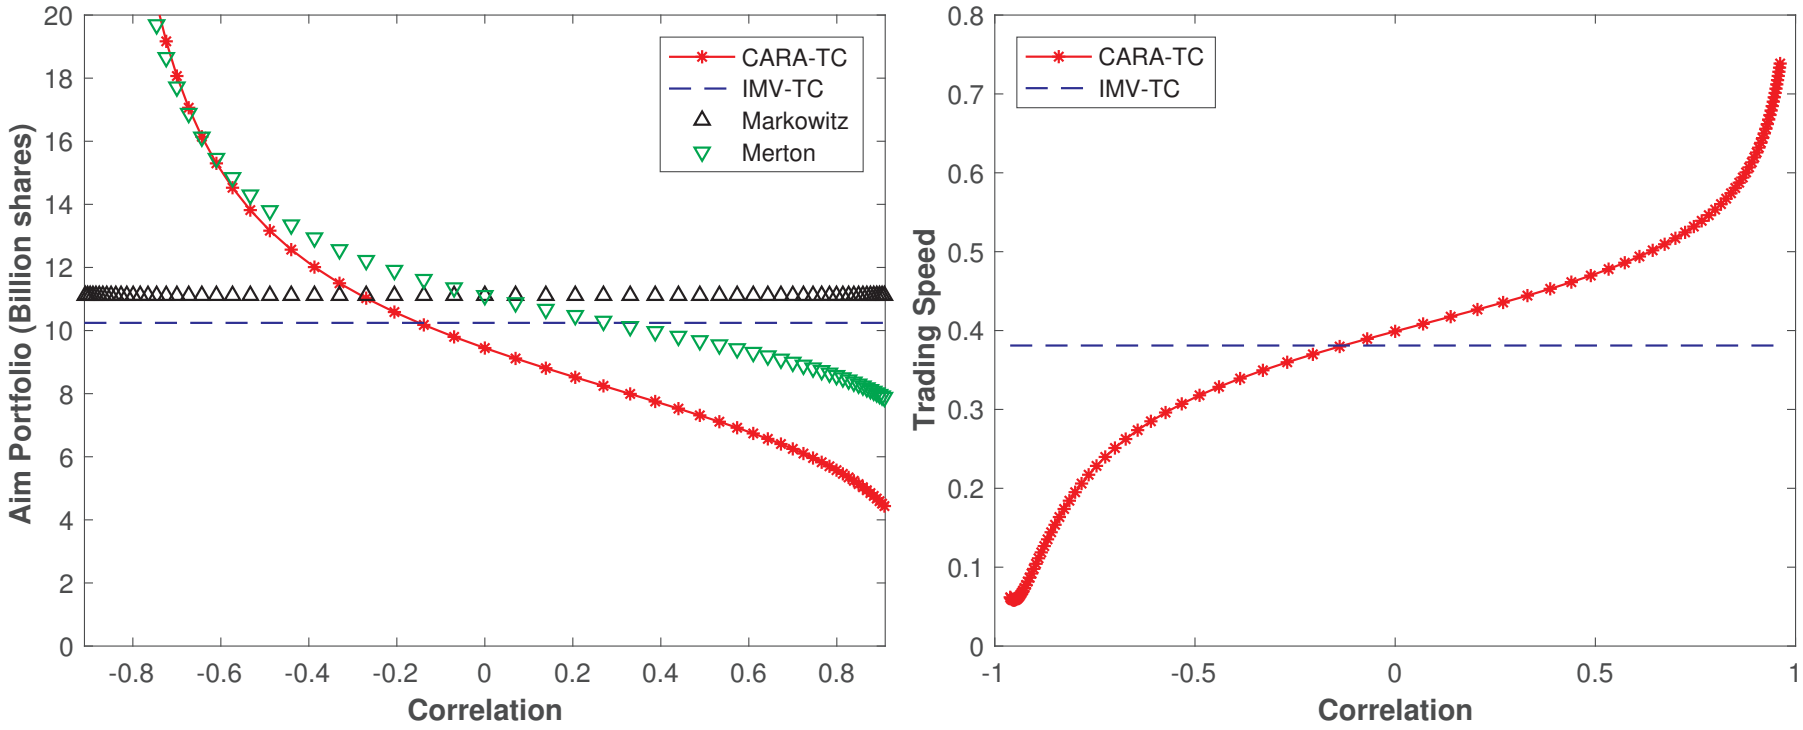

Figure: Parameters: $\mu_0 = 0$, $\mu = 1$, $\kappa = 0.1$, $\sigma_s = 0.3$, $\sigma_{xs} = -0.1$, $\sigma_x = 0.1$, $\Lambda = 2 \times 10^{-10}$, $\gamma_x = 10^{-9}$, $\gamma = 10^{-9}$, $x_0 = 1$, $\rho = 0.8$.

When signal is less persistent (high κ)

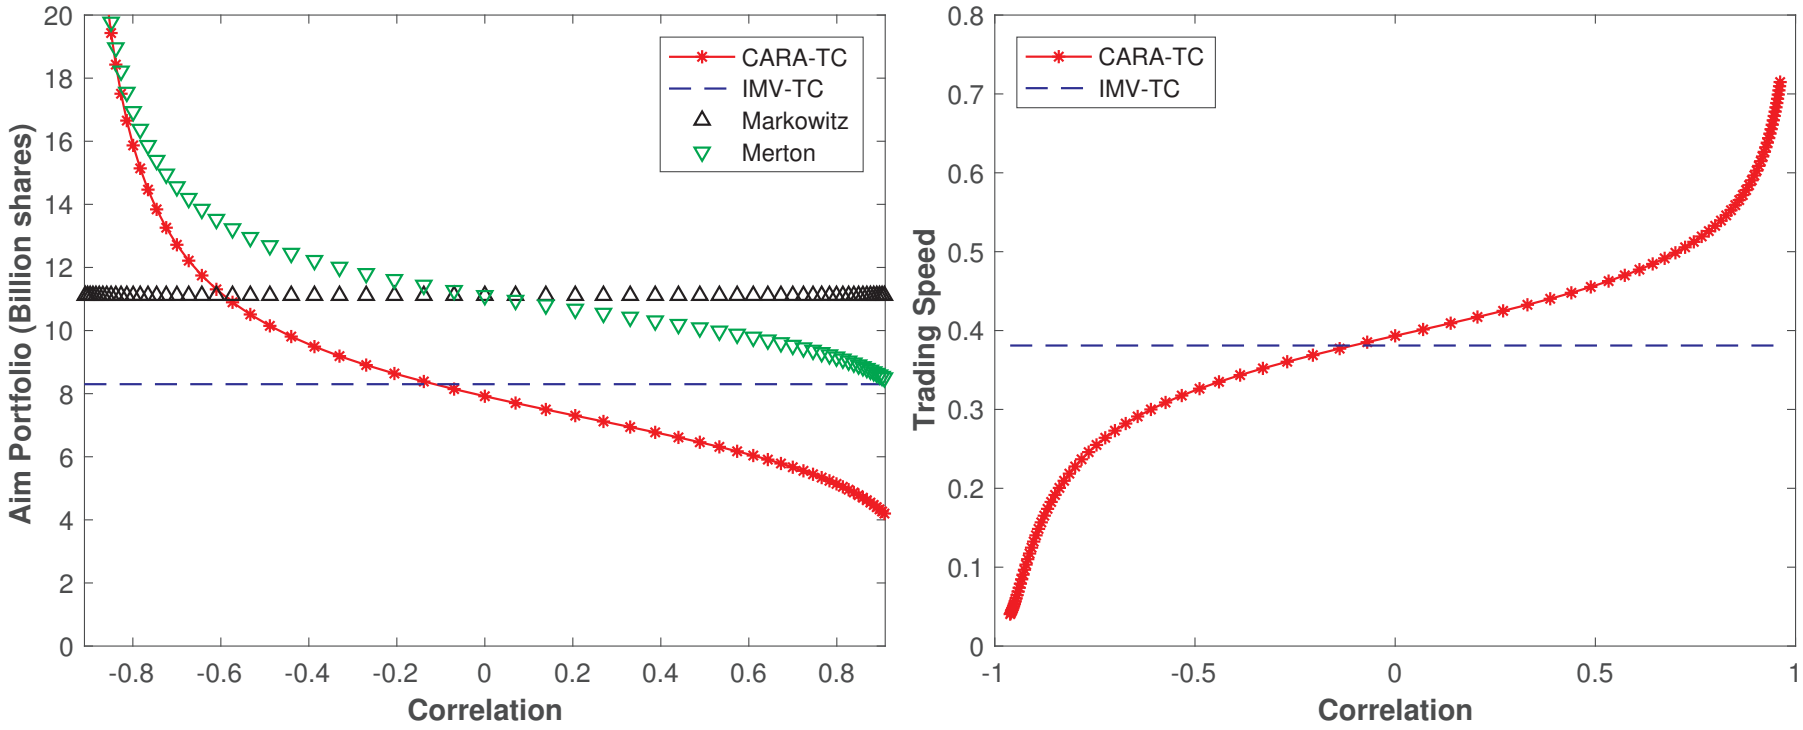

Figure: Parameters: $\mu_0 = 0$, $\mu = 1$, $\kappa = 0.4$, $\sigma_s = 0.3$, $\sigma_x = 0.1$, $\Lambda = 2 \times 10^{-10}$, $\gamma_x = 10^{-9}$, $\gamma = 10^{-9}$, $x_0 = 1$, $\rho = 0.8$.

When horizon is longer (low arrival intensity ρ)

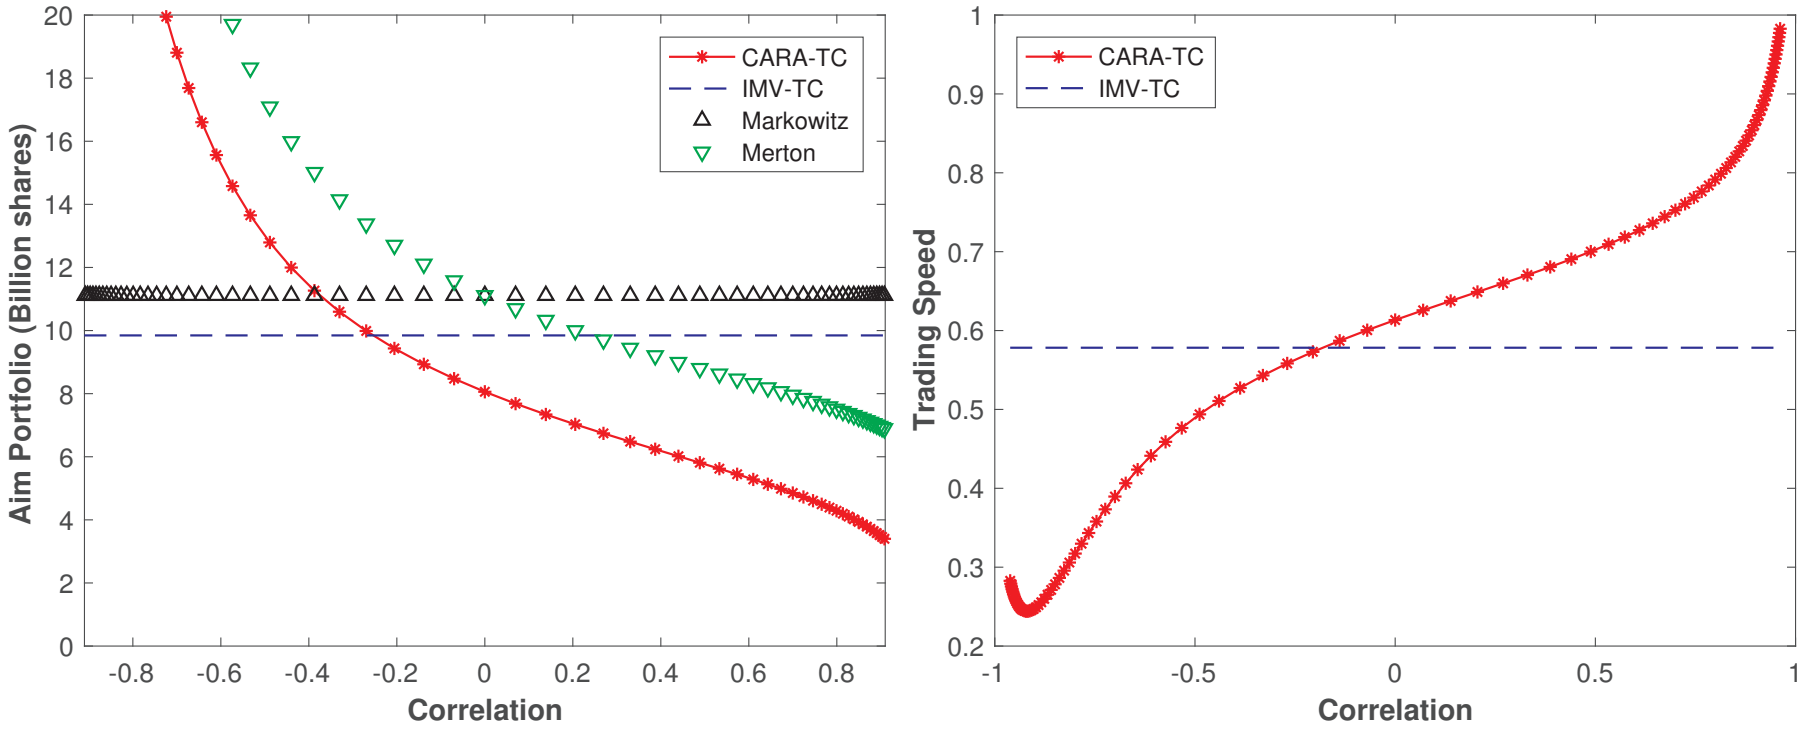

Figure: Parameters: $\mu_0 = 0$, $\mu = 1$, $\kappa = 0.1$, $\sigma_s = 0.3$, $\sigma_x = 0.1$, $\Lambda = 2 \times 10^{-10}$, $\gamma_x = 10^{-9}$, $\gamma = 10^{-9}$, $x_0 = 1$, $\rho = 0.2$.

Estimation of the return generating process: one stock-one latent factor

- Using Kalman filter, we estimate an unobservable state, x_t using time series of S&P 500 index (daily) price changes between 1962-07 and 2018-12:

$$dS_t = x_t dt + \sigma_s dZ_s(t) \tag{23}$$

$$dx_t = -\kappa x_t dt + \sigma_x dZ_x(t) + \sigma_{xs} dZ_s(t) \tag{24}$$

- Estimates using S&P 500 returns from 1962-07 to 2018-12:

Parameter	Value	Standard Error
κ	0.730	0.049
σ_s	0.137	0.008
σ_{xs}	-0.076	0.013
σ_x	0.005	0.029

→ Very negative correlation between dx_t and dS_t (< -0.95).

- For the experiments we set $\gamma = \gamma_x = 10^{-8}$ and $\rho = 0.05$.

Estimation of the transaction cost parameters

- Use proprietary institutional money managers Execution Data from the historical order databases of a large investment bank
- Use executions from top 50 stocks in terms of market capitalization
- Estimate Λ by fitting the following panel regression:

$$\Delta P_i = \Lambda \frac{D_i Q_i}{2} + \varepsilon_i$$

where

- Q_i is the number of shares and
- D_i is the direction of the trade.

	Dependent variable: $10^7 \Delta P$
Λ	1.1475***
	(0.217)
Note:	*p<0.1; **p<0.05; ***p<0.01

Trading Experiment

- Based on the filtered time series of x, the estimated return and t-cost parameters we:
  - Divide the sample into 112 trading intervals of 126 trading days (i.e., six-months).
  - At the beginning of each interval we start with \$0 and trade according to a trading rule (optimal, GP, myopic).
  - Record the stock position and the total accumulated wealth net of t-costs every day.
  - Compute the certainty equivalent, average across all intervals, for the CARA investor of various strategies.
- Results:

Statistic	Optimal	GP	Diff
CE	72,228	52,569	
Avg Util	-0.9993	-0.9995	1.96×10^{-4}
S.E	1.3×10^{-4}	0.9×10^{-4}	3.8×10^{-5}

- The aim portfolios as a function of x.

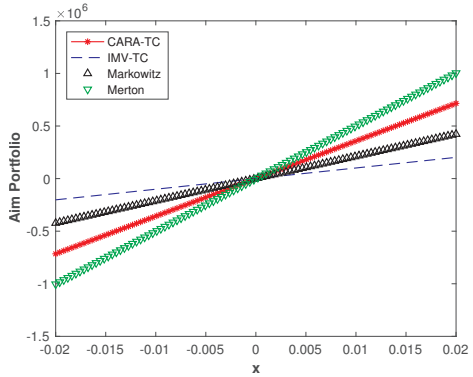

Average across all paths with TC

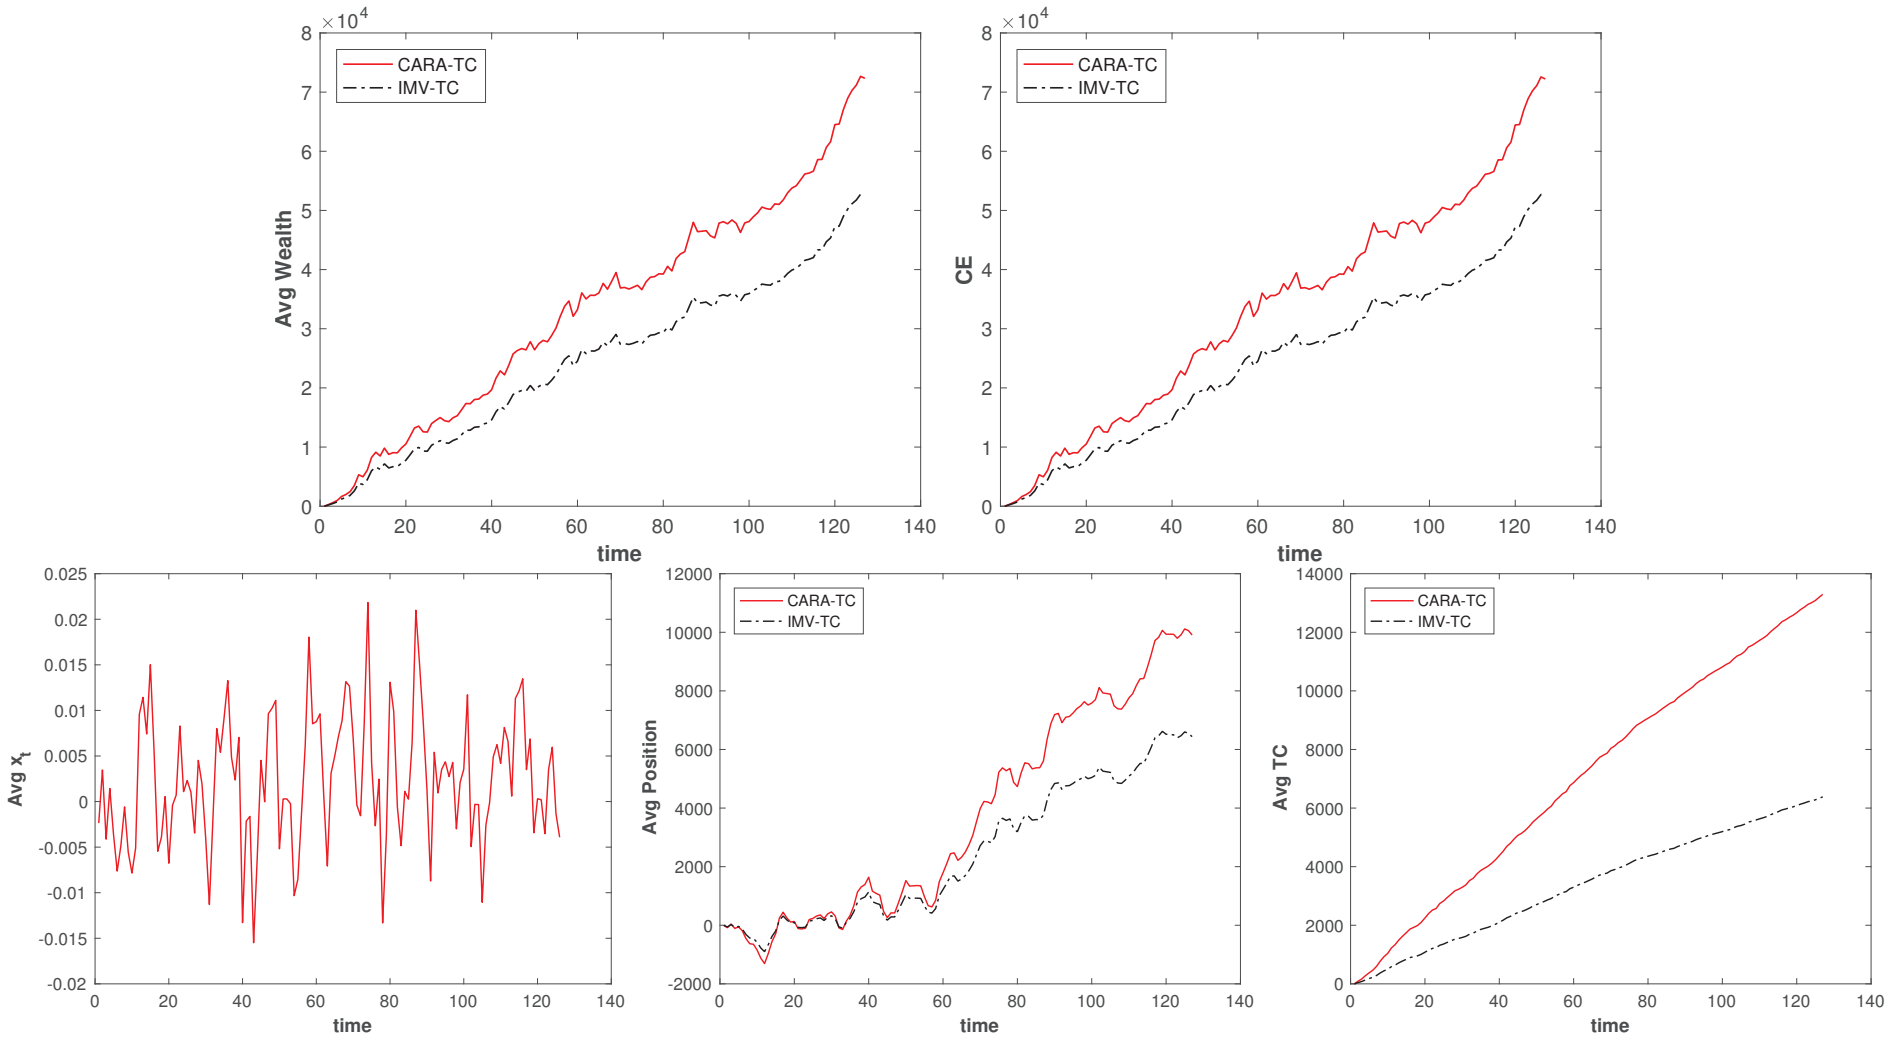

One Sample path with TC

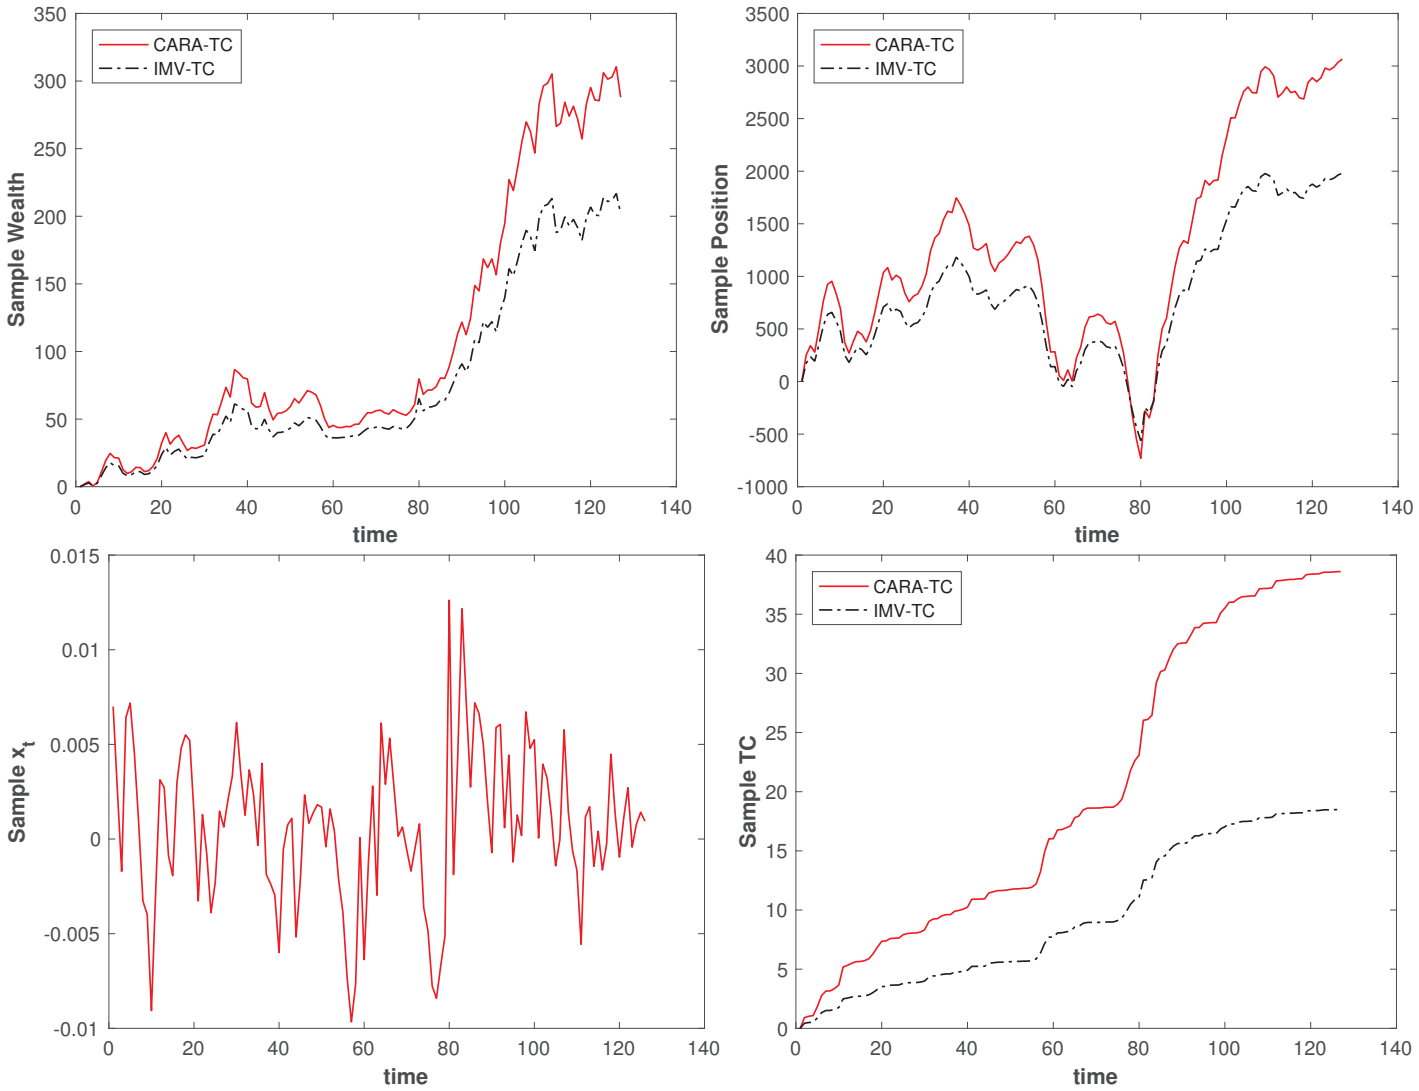

Average path without TC

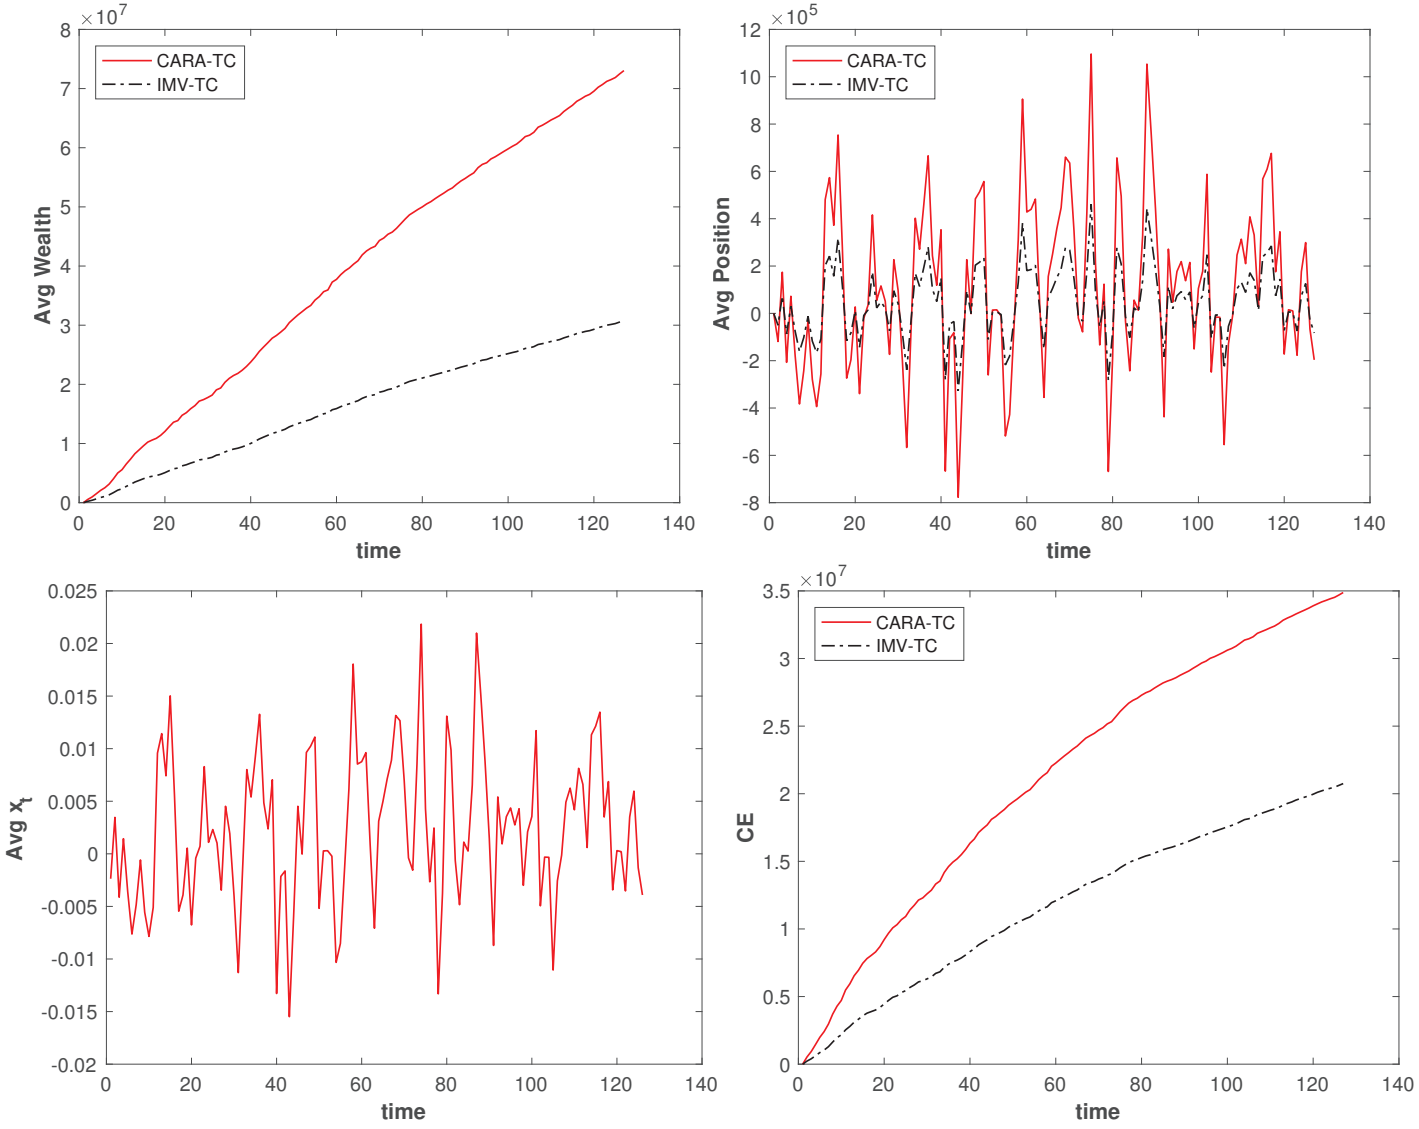

One sample path without TC

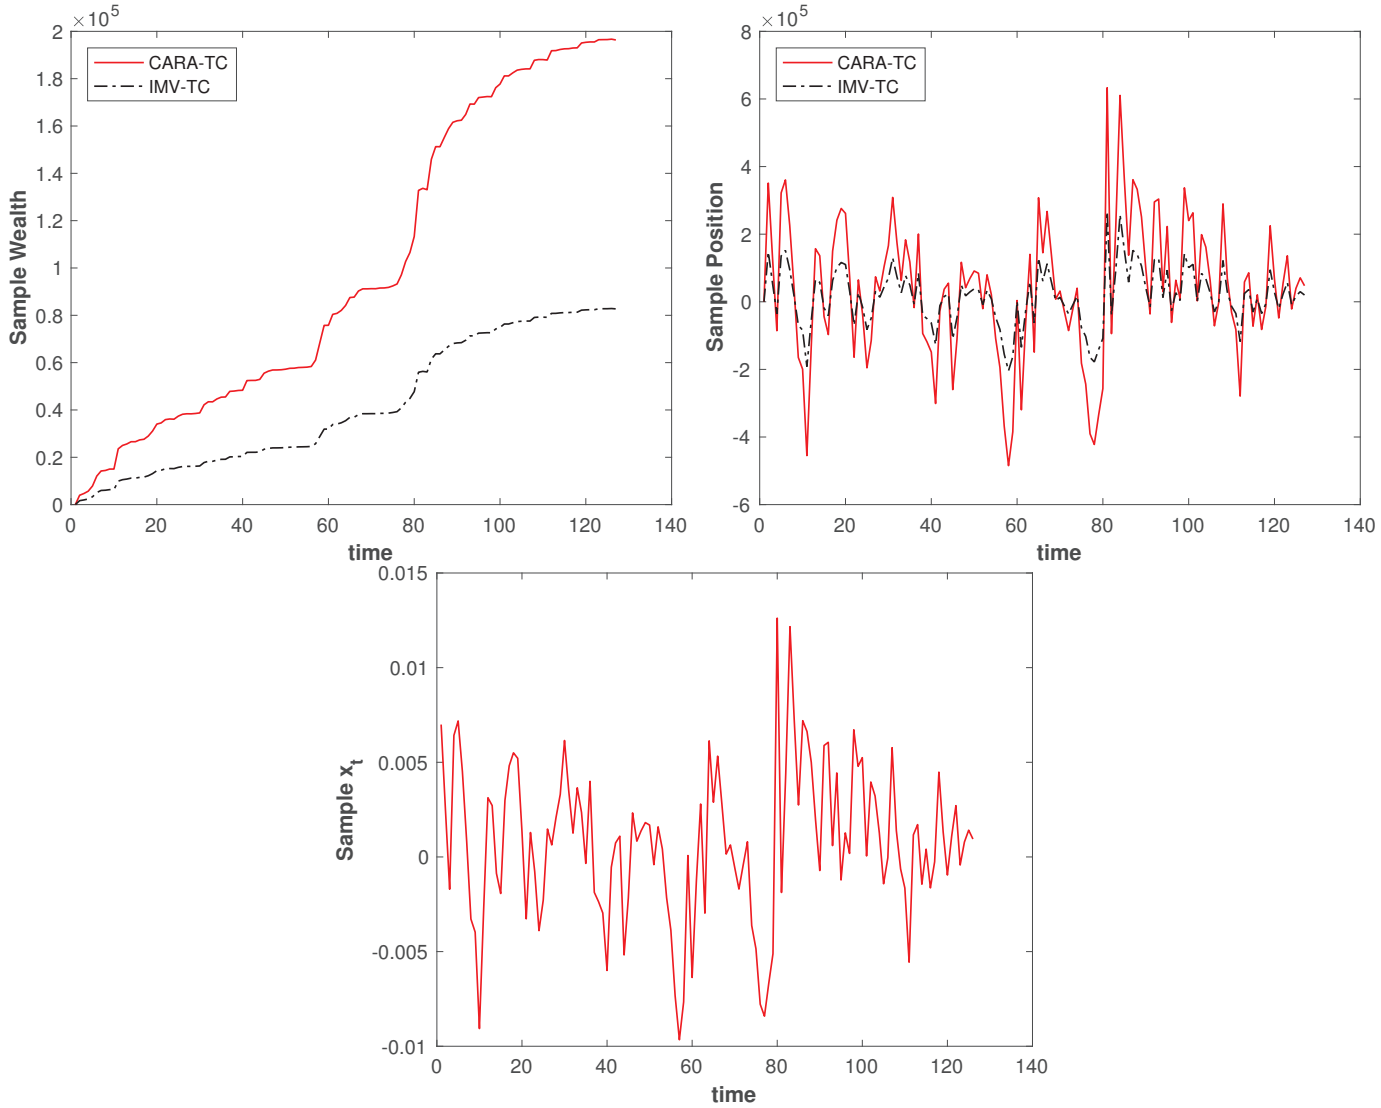

Conclusion

- We propose a set of preferences based on stochastic differential utility with source-dependent risk-aversion, which nest the widely used instantaneous mean-variance (Litterman (2005), Garleanu-Pedersen (2013)) and CARA expected utility.
- We derive an explicit solution for the portfolio choice problem in the presence of quadratic t-costs with arbitrary number of stocks and predictability in returns in terms of an optimal aim portfolio and trading speed.
- We show that, for a CARA investor, the hedging demand has large effect on optimal aim portfolio and trading speed, especially when the correlation between stock return and predictor is negative.
- In an in-sample experiment where we time the S&P 500 return based on its filtered latent predicted expected return, hedging demands significantly affect strategy performance.
- It remains to be seen whether this also holds out-of-sample.

Details on the execution data

- Data contains two trading algorithms:
  - volume weighted average price (VWAP) and
  - percentage of volume (PoV).
- Execution data covers S&P 500 stocks between January 2011 and December 2012.
- Execution duration is greater than 5 minutes but no longer than a full trading day.
- Total number of orders is 81,744 with an average size of approximately $1 million.
- The average participation rate of the order, the ratio of the order size to the total volume realized in the market, is approximately 6%.
- Back to main.