Optimal Dynamic Asset Allocation with Transaction Costs: The Role of Hedging Demands

Pierre Collin-Dufresne*
SFI@EPFL
email: pierre.collin-dufresne@epfl.ch

Kent Daniel
Columbia Business School and NBER
email: kd2371@columbia.edu

Mehmet Sağlam
University of Cincinnati
email: mehmet.saglam@uc.edu

This Draft: June 2023

Abstract

A number of papers have solved for the optimal dynamic portfolio strategy when expected returns are time-varying and trading is costly, but only for agents with myopic utility. Non-myopic agents benefit from hedging against shocks to the investment opportunity set even when transaction costs are zero (Merton, 1969, 1971). In this paper, we propose a solution to the dynamic portfolio allocation problem for non-myopic agents faced with a stochastic investment opportunity set, when trading is costly. We show that the agent's optimal policy is to trade toward an "aim" portfolio, the makeup of which depends both on transaction costs and on each asset's correlation with changes in the investment opportunity set. The speed at which the agent should trade towards the aim portfolio depends both on the shock's persistence and on the extent to which the shock can be effectively hedged.

JEL Classification: D53, G11, G12
Keywords: portfolio choice, dynamic models, transaction costs, hedging demand, price impact, mean-variance

*Please address correspondence to: Pierre Collin-Dufresne, Swiss Finance Institute at EPFL, Quartier UNIL-Dorigny, Extranef 209, CH-1015 Lausanne, Switzerland; Phone: +41 21 693 01 36; Email: pierre.collin-dufresne@epfl.ch

1 Introduction

Mean-variance efficient portfolio optimization, introduced by Markowitz (1952), is both a staple of MBA curricula and a critical tool for most quantitative asset managers. When either the vector of expected returns or the covariance matrix of returns is time-varying, a default solution is to simply hold the conditional mean-variance efficient 'Markowitz' (CMVE) portfolio. However, there are at least two reasons why it is not optimal for long-term investors to hold the CMVE portfolio: first, as shown in the seminal papers by Merton (1969, 1971) and Cox and Huang (1989) it may be optimal for long-term investors to deviate from the CMVE portfolio by tilting towards a portfolio whose realized returns are negatively correlated with changes in the CMVE portfolio's Sharpe ratio. Intuitively, holding this portfolio hedges the investor against changes in the investment opportunity set.

Second, if there are transaction costs then it may be too costly for investors to continuously and fully rebalance their portfolio in response to these shocks. Early papers (e.g., Constantinides, 1986; Davis and Norman, 1990; Dumas and Luciano, 1991) established that, with proportional transaction costs and with *i.i.d.* returns, it is optimal to refrain from trading until positions deviate substantially from the CMVE portfolio. More recently Litterman (2005) and Gârleanu and Pedersen (2013, GP) show that when expected returns are time-varying and price impact is linear (i.e., when transaction costs are quadratic), then it is optimal for investors to trade at a constant speed towards an *aim* portfolio, which puts less weight on stocks for which shocks to expected returns are less persistent.[1]

The latter set of papers obtain closed-form solutions for the optimal aim portfolio and trading speed, for arbitrary number of stocks and return forecasting factors, by relying on an ad-hoc conditionally mean-variance (CMV) objective function that leads to a standard linear-quadratic optimization problem, whose solution has been widely studied in mathematics and economics. Specifically, for an investor with wealth process W_t, the CMV objective is to maximize

$$(\star) \quad E\left[\int_0^\infty e^{-\rho t}\left\{dW_t - \frac{\gamma}{2}dW_t^2\right\}\right],$$

where γ can be interpreted as an instantaneous variance aversion coefficient. In the absence of transaction costs, this reduces to the myopic (instantaneous) mean-variance objective. Because it is very tractable in the presence of transaction costs or portfolio constraints, CMV has been widely used in the literature.[2]

In this paper, we propose an objective function which is equal to the certainty equivalent wealth of an agent with generalized recursive utility and with source-dependent constant absolute risk-aversion. Specifically, as in Skiadas (2008) and Hugonnier, Pelgrin, and St-Amour (2012), in

[1] Collin-Dufresne, Daniel, and Sağlam (2020) extend these results to a model where price impact and volatility are time-varying, and show how trading-speed and aim portfolio vary with volatility and transaction costs.

[2] In addition to the papers already cited, the CMV objective function is also used in Duffie and Zhu (2017); Du and Zhu (2017); Vayanos and Vila (2021); Gourinchas, Ray, and Vayanos (2021); Greenwood and Vayanos (2014); Malkhozov, Mueller, Vedolin, and Venter (2016); Danielsson, Shin, and Zigrand (2012).

our model the agent exhibits differential risk aversion to the shocks that drive price changes and changes in expected returns. We show that this preference specification is equally as tractable as the CMV framework. Moreover, these preferences converge to CMV preferences in the limit where the agent approaches risk-neutrality toward the risks driving expected returns and toward horizon risk for the stationary solution.[3] In the finite horizon case, when all risk-aversion coefficients towards all sources of risk are equal, this preference specification nests standard CARA (negative exponential) expected utility.

We characterize the closed-form solution to the optimal portfolio choice problem in a setting where the agent can trade a large number of securities whose expected returns are a linear function of a vector of mean-reverting state variables and where the agent faces quadratic trading costs. We show that the agent's optimal policy is to trade at a given rate towards an aim portfolio, where the aim portfolio is distinct from the optimal portfolio that the agent would choose in the absence of transaction costs. Since our framework nests both CARA expected utility and CMV preferences, we can investigate the impact of hedging demand on the aim portfolio and the trading speed.

We first show that, for a (myopic) agent with CMV preferences, the aim portfolio will be a trading-speed-discounted average of expected future CMVE portfolios, and the optimal trading speed matrix will be entirely determined by the ratio of stock volatility to price impact matrices. Thus, for a CMV-investor, the weight in the aim portfolio on a security with a mean-reverting expected return will always be lower than the weight in the corresponding CMVE portfolio and it will be *entirely independent* of the covariance matrix of signals. Specifically, consistent with the findings of GP, we show that CMV-investors underweight stocks with higher trading speed (that is, more volatile and more liquid stocks) and whose expected return decays faster. Further, CMV-investors aim portfolio and trading speed are identical whether signals are determinstic or stochastic.

However, when the agent is instead a long-term expected utility investor (ie., when their preferences deviate from CMV), we show that the preference to hedge changes in the investment opportunity set can dramatically change both the composition of the aim portfolio and speed at which the agent will trade towards the aim portolio. Both depend crucially on the correlation coefficient between the realized stock returns and the shocks to expected returns. If the correlation is sufficiently negative, we show that a long term CARA investor will choose to overweight a stock relative to the no-transaction-cost benchmark, despite it having positive price impact and mean-reverting return. GP's insight that, because of transaction costs, agents should underweight stocks with higher expected return decay rates, only holds *conditional* on a given level of correlation between signal and stock return. A long-term investor will optimally load more on a signal which mean reverts more quickly if that signal is also more negatively correlated with shocks to expected returns, i.e., if it better hedges changes in the investment opportunity set.[4]

[3] Specifically, we show that the infinite horizon objective ($\star$) above corresponds to the certainty equivalent of a source dependent recursive utility investor with a random horizon drawn from an exponential distribution with parameter ρ, who is risk-neutral towards expected return shocks and horizon risk.

[4] Note that in the absence of transaction costs, the optimal portfolio already holds a larger position than the

Furthermore, we find that the speed at which the agent will optimially trade towards the aim portolio depends on the correlation between shocks to expected returns and realized stock returns. In particular, the trading speed is decreasing for signals that are better hedges, i.e., for expected return shocks whose correlation with stock returns is more negative. As the correlation between realized returns and shocks to expected returns decreases, long-horizon stock return volatility decreases. Intuitively, the more negative this correlation, the better the ability of an asset to hedge future changes in expected return, and hence the lower the longer-horizon volatility. However, the presence of transaction costs increases the long-term investor's cost of hedging, and hence decreases the optimal trading speed. Finally, since the aim portfolio is a trading cost discounted value of future expected no-transaction-cost portfolios, negative correlation also implies a smaller discount between the no-transaction-cost optimum and the aim portfolio.

Related literature. Our paper is related to three strands of the dynamic portfolio choice literature. First, there is a large literature on the theory and the empirical relevance of hedging demand starting from Merton (1969, 1971). In particular, there are several studies examining how return predictability affects long-term asset allocation (see, among others, Brennan, Schwartz, and Lagnado, 1997; Brandt, 1999; Kim and Omberg, 1996; Campbell, 1999; Campbell and Viceira, 2002). In this literature, transaction costs are typically ignored, as the analytical solutions are typically not available in the presence of transaction costs.

Second, there are several academic papers studying the effect of transaction costs on dynamic portfolio choice but they typically focus on a very small number of assets (typically two) and limited use of return predictability (typically none). Constantinides (1986), Davis and Norman (1990), Dumas and Luciano (1991), Shreve and Soner (1994) study the two-asset (one risky and one risk-free) case with *i.i.d.* returns. Balduzzi and Lynch (1999) and Lynch and Balduzzi (2000) use a dynamic programming approach to investigate the impact of fixed and proportional transaction costs on the utility costs and the optimal rebalancing rule in a setting with a single risky asset with time-varying expected return. Longstaff (2001) studies a numerical solution in a setting with a single risky asset where this asset's returns have stochastic volatility, and when agents face liquidity constraints that force them to trade absolutely continuously. Liu (2004) studies the multi-asset case when agents have CARA preferences and when risky-asset returns are *i.i.d.*. Lynch and Tan (2010) use a numerical procedure to solve for the optimal portfolio choice of an investor with access to two risky assets under return predictability and proportional transaction costs. Brown and Smith (2011) discuss the high-dimensionality of the problem and provide approximately optimal trading strategies for a general dynamic portfolio optimization problem with transaction costs and return predictability that can be applied to larger number of stocks.

Third, there is a growing literature utilizing the tractability of the linear-quadratic formulation to derive closed-form solutions for the optimal investment portfolio in the presence of return predictability and transaction costs. Litterman (2005) and GP introduced this framework. They

CMV portfolio because of the positive hedging demand. The surprising finding is that transaction costs can actually further increase the overweighting of such stocks relative to the no-transaction-cost benchmark if the correlation is sufficiently negative.

demonstrate that it is optimal to trade away from the current portfolio and towards an "aim" portfolio which is a weighted average of the current and expected-future Markowitz portfolios on all future dates. Thus, the aim portfolio puts a higher weight on high expected return assets when that return is more persistent. In the GP setting, the speed at which the investor should move toward the aim portfolio is constant.

Collin-Dufresne, Daniel, and Sağlam (2020, CDS) consider a similar objective function (CMV utility with quadratic transaction costs) in a setting where expected returns, covariances and transaction costs are all stochastic. They find that the makeup of the aim portfolio and the trading speed are state-dependent, and vary with the relative magnitudes of transaction costs and state transition probabilities.[5]

2 The continuous time model with a finite horizon

Consider a continuous time economy where the N-dimensional vector of stock price processes S_t has dynamics:

$$dS_t = (\mu_0 + \mu x_t)dt + \sigma_s dZ_t^s \quad (1)$$

$$dx_t = -\kappa x_t dt + \sigma_x dZ_t^x + \sigma_{xs} dZ_t^s \quad (2)$$

We assume that the vector of expected return predictors x_t is K-dimensional and that the risk-free rate is zero.[6] Z^s and Z^x are vectors of independent Brownian motions that drive the randomness in stock prices and the state variables.[7] We define the instantaneous covariance matrix of returns to be Σ and the instantaneous covariance matrix of the innovations in the vector of state-variables to be Σ_x. Then, these covariance matrices are given by:

$$\Sigma = \sigma_s \sigma_s^\top, \quad (3)$$

$$\Sigma_x = \sigma_x \sigma_x^\top + \sigma_{xs} \sigma_{xs}^\top. \quad (4)$$

Remark 1 *Note that this specification nests the special case where each stock has an expected return driven by M stock specific predictors (e.g, book-to-market, momentum, reversal) that have different decay rates:*

$$dS_i(t) = (\mu_{0,i} + \sum_{m=1}^{M} \mu_{m,i} x_{m,i}(t))dt + \sigma_i dZ_i^s(t) \text{ for } i = 1, \ldots, N$$

$$dx_{j,i}(t) = -\kappa_j x_{j,i}(t)dt + \nu_{j,i} dZ_i^j(t) \text{ for } j = 1, \ldots, M.$$

[5] It would be interesting to extend our model to study how hedging demands driven by stochastic shifts in second moments affect their findings.

[6] For ease of reference and brevity, we will use 'returns' to refer to 'price changes' throughout the paper, consistent with many of the other papers in this literature.

[7] Since dZ_t^s is $N \times 1$, σ_x is $K \times K$, σ_s is $N \times N$ and σ_{xs} is $K \times N$.

To see this, set x to be the $(NM, 1)$ stacked vector of firm specific predictors and the matrix κ to be the (NM, NM) diagonal matrix whose diagonal coefficients cycle through the κ_m.[8]

The agent trades continuously by rebalancing the vector of number of shares n_t at an absolutely continuous rate θ_t, that is $dn_t = \theta_t dt$. When they rebalance they incur quadratic transaction costs so that their wealth process is given by:

$$dW_t = n_t^\top dS_t - \frac{1}{2}\theta_t^\top \Lambda \theta_t dt \tag{5}$$

$$= n_t^\top (\mu_0 + \mu\, x_t)dt + n_t^\top \sigma_s dZ_t^s - \frac{1}{2}\theta_t^\top \Lambda \theta_t dt \tag{6}$$

where Λ is a symmetric positive definite transaction-cost matrix.[9]

We assume that the agent maximizes her certainty equivalent wealth H_t, which is a process $(H_t, \sigma_{H,s}, \sigma_{H,x})$ which solves the following backward stochastic differential equation (BSDE):

$$H_t = \mathrm{E}_t \left[W_T - \int_t^T \left\{ \frac{1}{2}\gamma ||\sigma_{H,s}||^2 + \frac{1}{2}\gamma_x ||\sigma_{H,x}||^2 \right\} du \right] \tag{7}$$

Inspecting this equation we see that the solution H_t is the expected terminal wealth net of a risk-penalty, which is linear in the two components of its own variance that are due to the orthogonal Z^s and Z^x shocks, respectively. The agent attaches different 'source-specific' risk-aversion coefficients, γ and γ_x, to the two sources of risk, in the spirit of Skiadas (2008), and Hugonnier, Pelgrin, and St-Amour (2012). Our first result is to show that this certainty equivalent formulation nests two well-known objective functions: the constant absolute risk-aversion (CARA) expected utility and the conditional mean-variance (CMV) preferences.

Theorem 2 *The solution H_t to the recursive equation (7) is the certainty equivalent of an agent with source-dependent stochastic differential utility, who has a CARA coefficient γ towards Z^s shocks and γ_x towards Z^x shocks. It nests two important special cases:*

- *When $\gamma_x = \gamma$, it is the certainty equivalent of an agent with negative exponential CARA expected utility:*

$$H_t = -\frac{1}{\gamma} \log(\mathrm{E}_t[e^{-\gamma W_T}]). \tag{8}$$

- *When $\gamma_x \sigma_x = 0$ and $\sigma_{xs} = 0$, it reduces to the CMV objective function:*

$$H_t = W_t + \mathrm{E}_t \left[\int_t^T \left\{ dW_u - \frac{1}{2}\gamma dW_u^2 \right\} \right]. \tag{9}$$

Proof. See Appendix A and Appendix B. ■

[8] Other matrices need to be adjusted appropriately as well. For example, μ is the (N, NM) diagonal sparse matrix which has row vector $[\mu_{1,i}, \mu_{2,i}, \ldots, \mu_{N,i}]$ on the i^{th} 'diagonal.'

[9] Assuming Λ is positive definite insures that transaction costs on any non-zero trade must be strictly positive. Assuming it is symmetric is without loss of generality given the quadratic form of the transaction costs.

This theorem shows that the certainty equivalent H_t defined in equation (7) nests both CARA and CMV preferences. Because of its analytical tractability, the CMV framework has been widely used in the literature on dynamic portfolio choice with transaction costs (e.g., Litterman, 2005; Gârleanu and Pedersen, 2013; Collin-Dufresne, Daniel, and Sağlam, 2020), with holding costs (e.g., Duffie and Zhu, 2017) and with portfolio constraints (e.g., Vayanos and Vila, 2021). The second result of the theorem show, that when expected returns are non stochastic (i.e., when $\sigma_x = \sigma_{xs} = 0$), then the optimal portfolio for CARA and CMV investors is identical. However, when the expected returns are stochastic, the solutions diverge. In this latter setting, we can demonstrate the following:

Corollary 3 *The CMV objective function of equation (9) reduces to the linear-quadratic framework used in Litterman (2005), Gârleanu and Pedersen (2013), and Collin-Dufresne, Daniel, and Sağlam (2020):*

$$J_t := H_t - W_t = \mathrm{E}_t \left[\int_t^T \left\{ n_u^\top (\mu_0 + \mu\, x_u) du - \frac{1}{2} \theta_u^\top \Lambda \theta_u - \frac{1}{2} \gamma n_u^\top \Sigma n_u \right\} du \right] \quad s.t.\ dn_t = \theta_t dt. \quad (10)$$

Its solution is identical to that of an agent with source dependent utility who maximizes the certainty equivalent (7), is risk-neutral to state-variable shocks (i.e., $\gamma_x = 0$), uses the correct covariance matrix for both stock returns (Σ) and state variables (Σ_x), but assumes zero correlation between the two, that is $\sigma_{xs} = 0$.

In the absence of transaction costs, it is optimal for the CMV agent to act myopically and continuously rebalance towards the conditional mean-variance efficient (CMVE) portfolio. However, even in the absence of transaction costs, the CARA investor optimally deviates from the CMVE portfolio in order to hedge shocks to the investment opportunity set (Merton, 1971).

When transaction costs are non-zero, Gârleanu and Pedersen (2013) show that it is optimal for the CMV-investor to trade at a constant rate towards an *aim*-portfolio, that can be interpreted as a discounted average of expected future CMVE portfolios (note that CMVE portfolios vary stochastically as the expected returns are driven by x_t).[10]

Our contribution is to consider the optimal dynamic portfolio for an agent with long-horizon preferences (e.g., a CARA investor) in a setting with a stochastic investment opportunities, and where transaction costs are non-zero. Specifically, we characterize the optimal trading strategy of the source-dependent utility agent (which nests both CMV and CARA) in the presence of transaction costs. We would like to understand whether and how the seminal insight of Merton (1971)—that a long-term investor should deviate from her myopic portfolio to take advantage of stock predictability—is affected by the presence of transaction costs. Is it still possible to characterize the optimal trading strategy of a non-myopic agent in terms of an aim-portfolio and trading speed, as in GP? How do hedging demands affect the aim portfolio and trading speed?

The following theorem describes the solution to the optimal portfolio choice problem of the agent with recursive utility with source-dependent risk-aversion.

[10] See Theorem 5 below for a precise restatement of this result in the context of our model.

Theorem 4 *Suppose an agent maximizes her certainty equivalent H_t defined in equation (7) by choosing her optimal position vector n_t given wealth dynamics described by equation (6).*

If there are no transaction costs ($\Lambda = 0$), *then the maximum certainty equivalent is $H_t = W_t + J(x_t, t)$ where*

$$J(x,t) = c_0(t) + c_1(t) + \frac{1}{2}x^\top c_2(t)x, \tag{11}$$

where the (matrix) functions c_1, c_2 solve the system of ODEs:

$$-\dot{c}_1 = (\mu - \gamma\Sigma_{sx}c_2)^\top (\gamma\Sigma)^{-1}\mu_0 - \{(\mu - \gamma\Sigma_{sx}c_2)^\top\Sigma^{-1}\Sigma_{sx} + c_2^\top\Omega + \kappa^\top\}c_1 \tag{12}$$

$$-\dot{c}_2 = c_2^\top \left(\gamma\Sigma_{sx}^\top\Sigma^{-1}\Sigma_{sx} - \Omega\right) c_2 + 2c_2^\top(-\kappa - \Sigma_{sx}^\top\Sigma^{-1}\mu) + \mu^\top(\gamma\Sigma)^{-1}\mu \tag{13}$$

where

$$\Omega = \gamma\sigma_{xs}\sigma_{xs}^\top + \gamma_x\sigma_x\sigma_x^\top, \tag{14}$$

$$\Sigma_{sx} = \sigma_s\sigma_{xs}^\top, \tag{15}$$

and the boundary conditions are given by $c_1(T) = 0$ and $c_2(T) = 0$. c_0 is given in equation (106) in Appendix E. The optimal position (in the absence of transaction costs) is given by:

$$n_t = (\gamma\Sigma)^{-1}(\mu_0 + \mu x_t) - \Sigma^{-1}\Sigma_{sx}(c_1(t) + c_2(t)x) \tag{16}$$

In particular, if $\Sigma_{sx} = 0$ then it is optimal to hold the CMVE Markowitz portfolio:

$$CMVE_t = (\gamma\Sigma)^{-1}(\mu_0 + \mu x_t). \tag{17}$$

If Λ is positive definite, *then the maximum certainty equivalent is $H_t = W_t + J(n_t, x_t, t)$ where*

$$J(n,x,t) = -\frac{1}{2}n^\top Q(t)n + n^\top(q_0(t) + q(t)^\top x) + c_0(t) + c_1(t)^\top x + \frac{1}{2}x^\top c_2(t)x, \tag{18}$$

where the (matrix) functions Q, q, q_0, c_1, c_2 solve the system of ODEs:[11]

$$-\dot{Q} = \gamma\Sigma - Q\Lambda^{-1}Q + q^\top\Omega q + \gamma(\Sigma_{sx}q + q^\top\Sigma_{sx}^\top) \tag{19}$$

$$-\dot{q}^\top = \mu - q^\top\kappa - Q\Lambda^{-1}q^\top - q^\top\Omega c_2 - \gamma\Sigma_{sx}c_2 \tag{20}$$

$$-\dot{c}_2 = -(c_2\kappa + \kappa^\top c_2) + q\Lambda^{-1}q^\top - c_2\Omega c_2 \tag{21}$$

$$-\dot{q}_0 = \mu_0 - Q\Lambda^{-1}q_0 - q^\top\Omega c_1 - \gamma\Sigma_{sx}c_1 \tag{22}$$

$$-\dot{c}_1 = -\kappa^\top c_1 + q\Lambda^{-1}q_0 - c_2\Omega c_1 \tag{23}$$

[11]Below we write the system of ODE for the symmetric Q, c_2 matrices, which is without loss of generality given the quadratic objective function.

subject to boundary conditions $Q(T) = 0$*,* $q(T) = 0$*,* $q_0(T) = 0$*,* $c_1(T) = 0$ *and* $c_2(T) = 0$*.* c_0 *is given in equation (119) in Appendix F.*

The optimal trading strategy is to trade at a deterministic (matrix valued) trading rate τ_t *towards an optimal aim portfolio such that:*

$$dn_t = \tau_t(aim(x_t, t) - n_t)\, dt \tag{24}$$

$$\tau_t = \Lambda^{-1}Q(t) \tag{25}$$

$$aim(x_t, t) = Q(t)^{-1}(q_0(t) + q(t)^\top x_t) \tag{26}$$

We note that the optimal aim portfolio corresponds to the position that maximizes the value function, that is $aim(t, x) = \operatorname{argmax}_n J(n, x, t)$*.*

Proof. The derivation of the solution without transaction costs ($\Lambda = 0$) is in Appendix E. The proof of the case with transaction costs is in Appendix F. ■

The optimal trading strategy for the agent with source dependent utility—summarized in equations (24)–(26)—takes a form similar to the solutions identified in GP or CDS: the strategy moves away from the current portfolio n_t towards an aim portfolio $aim(x_t, t)$ at a rate of τ_t. However, in contrast with the GP solution, the aim portfolio now depends on the state vector x_t, while the trading speed remains deterministic.

This solution nests these other findings, in that:

1. Consistent with Merton (1971), even when transaction costs are zero, if security returns are correlated with innovations in the state-variables (ie., when $\sigma_{xs} \neq 0$), then the agent will hold a portfolio that differs from the CMVE Markowitz portfolio. This is true unless the agent has CMV preferences, in which case $c_1(t)$ and $c_2(t)$ will be zero.

2. Consistent with GP, when transaction costs are non-zero, it will be optimal to deviate from the CMVE Markowitz portfolio even when $\sigma_{xs} = 0$.

3. The optimal trading strategy of GP is independent of covariance matrix of the state-variable. That is the investor makes the same investment decision if x_t were a deterministic process (even though her expected utility level is affected by state variable volatility)!

In our generalized setting, there are at least two reasons why the aim-portfolio will deviate from the GP/CDS solution in which the aim portfolio is a weighted average of expected future MVE portfolios: a traditional "Merton" no-transaction-cost investment-opportunity-set-hedging demand, and a transaction-cost specific hedging demand.

To understand both components, we next give a few analytical results that characterize the solution to the CMV objective function (which corresponds to the case where $\gamma_x = 0$ and $\sigma_{xs} = 0$).In this case the system has a closed-form solution that is similar to that obtained in GP. It can be characterized fully in terms of the eigenvalue decomposition of the matrix $\gamma\Lambda^{-1}\Sigma$. Specifically, we define (η, F) to be the vector of eigenvalues and the matrix of eigenvectors so that

$$\gamma\Lambda^{-1}\Sigma = FD_\eta F^{-1} \tag{27}$$

where D_η is the diagonal matrix with eigenvalue η_i on the i^{th} diagonal, and F is corresponding matrix of eigenvectors. Then we have the following result:

Theorem 5 *When $\gamma_x = 0$ and $\sigma_{xs} = 0$, the optimal trading speed matrix, $\tau_t = \Lambda^{-1}Q(t)$, is given by:*

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

The optimal aim portfolio of the investor with CMV preferences $aim(x,t) = Q(t)^{-1}(q_0(t) + q(t)^\top x)$ can be interpreted as a Markowitz portfolio where we replace the expected return vector by a trading-speed weighted average of future expected returns:

$$aim(x,t) = (\gamma\Sigma)^{-1}\int_t^T \omega_{t,u}\mu_S(t,u)du \tag{28}$$

$$\omega_{t,u} = (\int_t^T e^{-\int_t^z \tau_s^\top ds}dz)^{-1}e^{-\int_t^u \tau_s^\top ds} \tag{29}$$

where we define the expected future stock return by

$$\mu_S(t,u) = \frac{1}{dt}\mathrm{E}_t[dS_u] = \mu_0 + \mu e^{-\int_t^u \kappa ds}x_t \tag{30}$$

The CMV-agent portfolio is independent *of the covariance matrix (σ_x, σ_{xs}) of the expected return.*

Proof. The proof is provided in Appendix G. ■

We observe that the optimal aim portfolio of the investor with CMV-preferences has the same form as the Markowitz portfolios, but where the loadings μ on the time-varying return predictors, x_t, are modified to account for the combination of (i) transaction costs ($\omega_{t,u}$) and (ii) persistence (κ weights). Note that the ω weights only depend on the trading speed τ_t. Further, they are strictly positive and integrate to one, that is $\int_t^T \omega_{t,u}du = 1$. This can be interpreted as an 'average trade horizon': the higher the trading speed is, the shorter the horizon and the more we discount the future expected factor returns. In addition, since factors with higher κ are expected to revert faster towards zero,[12] the solution implies we should also underweight more, relative to the Markowitz portfolio, factors which are less persistent (i.e., with a higher mean-reversion rate κ). In particular, if factors are driven only by permanent shocks, that is $\kappa = 0$, then the optimal aim portfolio is the Markowitz portfolio (since the ω-weights integrate to one by construction).

[12] Recall that $\mathrm{E}_t[x_u] = e^{-\int_t^u \kappa ds}x_t$.

For now we have worked in a finite horizon setting where the link between the CARA normal setting and the instantaneous mean-variance framework used in the literature is the most straightforward to demonstrate. To avoid the explicit time-dependence introduced by the finite horizon setting, it is useful to extend the setting to an infinite horizon discounted objective function. This is also the choice made in GP, and CDS. We next show how to generalize this section's results to a stationary objective function with infinite horizon and demonstrate the connection to the certainty equivalent wealth of a source dependent risk-aversion agent with a random horizon.

3 The stationary model with a random horizon

It is natural to consider the stationary problem where we assume that the horizon $\mathcal{T}$ is drawn from an exponential distribution with parameter ρ. In that case we assume that the agent maximizes her certainty equivalent which is a process $(H_t, \sigma_{H,s}, \sigma_{H,x})$ which solves the following backward stochastic differential equation (BSDE):

$$H_t = \mathrm{E}_t\left[W_{\mathcal{T}} - \int_t^{\mathcal{T}} \left\{\frac{1}{2}\gamma||\sigma_{H,s}||^2 + \frac{1}{2}\gamma_x||\sigma_{H,x}||^2\right\} du\right] \tag{31}$$

$$= W_t + \mathrm{E}_t\left[\int_t^{\infty} e^{-\rho(u-t)}(dW_u - \left\{\frac{1}{2}\gamma||\sigma_{H,s}||^2 + \frac{1}{2}\gamma_x||\sigma_{H,x}||^2\right\} du)\right] \tag{32}$$

One might think that this stationary version of equation (7) should correspond to the certainty equivalent of a CARA agent who maximizes $\mathrm{E}[-e^{-\gamma W_{\mathcal{T}}}]$ for $\gamma_x = \gamma$. However, we show in the following theorem that this is not the case. Instead, the objective function (32) corresponds to that of an agent with source dependent risk-aversion who is risk-neutral with respect to horizon risk. When we add the risk of a random horizon arrival $\mathcal{T}$ to the Brownian risks, (Z^s, Z^x), the CARA agent is also risk-averse to that new source of risk and requires an extra premium, as we illustrate in Remark 7 below. As we show in the next theorem, the objective function in (31)- (32) corresponds to an agent who does not require a premium for horizon risk. The following theorem makes this explicit.

Theorem 6 *On the filtered probability space generated by $(Z^s, Z^x, \mathbf{1}_{\{\mathcal{T}\le t\}})$, consider the process $(H_t, \sigma_{H,s}, \sigma_{H,x})$ which solves the following backward stochastic differential equation (BSDE):*

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

Then H_t is the certainty equivalent of an agent with source-dependent constant absolute risk-aversion, with CARA γ toward Z^s shocks, γ_x towards Z^x shocks, and $\gamma_{\mathcal{T}}$ towards the horizon arrival shock, $\mathbf{1}_{\{\mathcal{T}\le t\}}$, which triggers a jump in H. It nests the special cases:

- *When $\gamma_{\mathcal{T}} = \gamma_x = \gamma$, it is the certainty equivalent of an agent with negative exponential CARA* *expected utility:*

$$H_t = -\frac{1}{\gamma}\log(\mathrm{E}_t[e^{-\gamma W_{\mathcal{T}}}]). \tag{33}$$

- *When* $\gamma_{\mathcal{T}} = 0$, *it reduces to the objective function (a stationary version of (7)) proposed in (32).*
- *When* $\gamma_{\mathcal{T}} = 0$, $\gamma_x\sigma_x = 0$ *and* $\sigma_{xs} = 0$, *it reduces to the discounted CMV objective function:*

$$H_t = W_t + \mathrm{E}_t\left[\int_t^\infty e^{-\rho(u-t)}\left\{dW_u - \frac{1}{2}\gamma dW_u^2\right\}\right]. \tag{34}$$

Proof. The proof is provided in Appendix D. ■

Remark 7 *To understand why a CARA investor dislikes horizon risk, consider the simple case where* $dW_t = \mu dt + \sigma dZ_t^s$, *that is wealth is solely driven by one Brownian motion. Then, consider the expected utility of the CARA agent*

$$E[-e^{-\gamma W_{\mathcal{T}}}] = -e^{-\gamma W_0}\int_0^\infty \rho e^{-\rho t-\gamma(\mu-\frac{1}{2}\gamma\sigma^2)t}dt = -\frac{e^{-\gamma W_0}}{1+\frac{\gamma(\mu-\frac{1}{2}\gamma\sigma^2)}{\rho}}.$$

Her expected utility of terminal wealth at the expected arrival time $\mathrm{E}[\mathcal{T}] = \frac{1}{\rho}$ *is given by:*

$$\mathrm{E}[-e^{-\gamma W_{1/\rho}}] = -e^{\frac{-\gamma(W_0+\mu-\frac{1}{2}\gamma\sigma^2)}{\rho}}$$

Since $e^z > 1+z$ *for all* $z \neq 0$ *and in particular for* $z = \frac{\gamma(\mu-\frac{1}{2}\gamma\sigma^2)}{\rho}$ *we see that for this CARA agent:*

$$\mathrm{E}[U(W_{\mathcal{T}})] < \mathrm{E}[U(W_{\mathrm{E}[\mathcal{T}]})] \quad \Longleftrightarrow \quad \frac{\gamma(\mu-\frac{1}{2}\gamma\sigma^2)}{\rho} \neq 0$$

This follows from Jensen's inequality. We see that a risk-premium for horizon risk arises as soon as the expected return on total wealth does not exactly compensate the investor for its diffusion risk (in the example as long as $\mu - \frac{1}{2}\gamma\sigma^2 \neq 0$*). If the agent's terminal wealth were guaranteed and independent of the horizon (i.e.,* $\mu = \sigma = 0$ *in the example) then, a consequence of time-separable utility, is that the agent would not care about horizon risk. With CARA utility the risk-aversion coefficient associated with the horizon risk* $\mathcal{T}$ *is the same as that associated to the Brownian motion shocks* Z^s, Z^x *that drive financial wealth. Instead, with our source-dependent utility, the agent can have different risk-aversion coefficients associated with the three different sources of risk. The standard discounted CMV preferences used in GP, Litterman, and others correspond to an agent who is risk-neutral towards horizon risk.*

In the following we focus on the solution of the agent with preferences given in (32), which corresponds to the stationary version of the problem considered in the previous section. The following theorem describes the optimal solution, and is the analogue to Theorem 4 with an infinite horizon.

Theorem 8 *Suppose an agent maximizes her certainty equivalent H_t defined in equation (32) by choosing her optimal position vector n_t given wealth dynamics given in equation (6).*

If there are no transaction costs ($\Lambda = 0$), *then the maximum certainty equivalent is $H_t = W_t + J(x_t)$ where*

$$J(x) = c_0 + c_1^\top x + \frac{1}{2} x^\top c_2 x, \tag{35}$$

where the (matrix) functions c_1, c_2 solve the system of ODEs:

$$\rho c_1 = (\mu - \gamma \Sigma_{sx} c_2)^\top (\gamma\Sigma)^{-1}\mu_0 - \{(\mu - \gamma\Sigma_{sx}c_2)^\top \Sigma^{-1}\Sigma_{sx} + c_2^\top \Omega + \kappa^\top\} c_1 \tag{36}$$

$$\rho c_2 = c_2^\top \left(\gamma \Sigma_{sx}^\top \Sigma^{-1}\Sigma_{sx} - \Omega\right) c_2 + 2c_2^\top(-\kappa - \Sigma_{sx}^\top\Sigma^{-1}\mu) + \mu^\top(\gamma\Sigma)^{-1}\mu. \tag{37}$$

The equation for c_0 is provided in the Appendix. The optimal position (in the absence of transaction costs) is:

$$n_t = (\gamma\Sigma)^{-1}(\mu_0 + \mu x_t) - \Sigma^{-1}\Sigma_{sx}(c_1 + c_2 x) \tag{38}$$

Note that, in particular, if $\Sigma_{sx} = 0$, then it is optimal to hold the CMVE Markowitz portfolio.

If Λ is positive definite, *the maximum certainty equivalent is $H_t = W_t + J(n_t, x_t)$, where*

$$J(n, x) = -\frac{1}{2} n^\top Q n + n^\top(q_0 + q^\top x) + c_0 + c_1^\top x + \frac{1}{2}x^\top c_2 x, \tag{39}$$

where the coefficient matrices Q, q, q_0, c_1, c_2 solve the system of ODEs:[13]

$$\rho Q = \gamma\Sigma - Q\Lambda^{-1}Q + q^\top \Omega q + \gamma(\Sigma_{sx} q + q^\top \Sigma_{sx}^\top) \tag{40}$$

$$\rho q^\top = \mu - q^\top\kappa - Q\Lambda^{-1}q^\top - q^\top\Omega c_2 - \gamma\Sigma_{sx}c_2 \tag{41}$$

$$\rho c_2 = -(c_2\kappa + \kappa^\top c_2) + q\Lambda^{-1}q^\top - c_2\Omega c_2 \tag{42}$$

$$\rho q_0 = \mu_0 - Q\Lambda^{-1}q_0 - q^\top\Omega c_1 - \gamma\Sigma_{sx}c_1 \tag{43}$$

$$\rho c_1 = -\kappa^\top c_1 + q\Lambda^{-1}q_0 - c_2\Omega c_1 \tag{44}$$

and c_0 is given in the Appendix.

The optimal trading strategy is to trade at a stock-specific constant trading rate (matrix) τ towards an optimal aim portfolio such that:

[13] Below we write the ODE for symmetric Q, c_2 matrices, wlog.

$$dn_t = \tau(aim(x_t) - n_t)\, dt \quad (45)$$
$$\tau = \Lambda^{-1} Q \quad (46)$$
$$aim(x) = Q^{-1}(q_0 + q^\top x) \quad (47)$$

We note that the optimal aim portfolio corresponds to the position that maximizes the value function, that is $aim(x) = \text{argmax}_n J(n, x)$.

Proof. The derivation of this solution (with $\Lambda = 0$) is given in appendix H. The derivation of the solution of the case with non-zero transaction costs is given in appendix I. ■

Thus, as in the finite horizon case described in the previous section, the optimal trading strategy for the agent with source dependent utility has the same form as that obtained in GP or CDS. Specifically, it is optimal to trade from the current position n_t towards an aim portfolio $aim(x_t)$ at a constant trading speed matrix τ.

To better understand the role of hedging demands in shaping the aim portfolio, we will compare numerically in the following section the optimal solution for the CARA agent to that of the CMV investor. Recall that in he absence of transaction costs, the CMV investor always holds the CMVE portfolio. With transaction costs however, the solutin of the CMV investor can be characterized explicitly (setting $\gamma_x = 0$ and $\sigma_{xs} = 0$ in theorem 5), in terms of the eigen-value decomposition of the matrix $\gamma\Lambda^{-1}\Sigma$. Specifically, we define (η, F) to be the vector of eigenvalues and the matrix of eigenfactors so that

$$\gamma\Lambda^{-1}\Sigma = F D_\eta F^{-1} \quad (48)$$

where D_η is the diagonal matrix with eigenvalue η_i on the i^{th} diagonal. Then we have the following result:

Theorem 9 *When* $\gamma_x = 0$ *and* $\sigma_{xs} = 0$ *then the optimal trading speed matrix* $\tau = \Lambda^{-1}Q$ *is given by:*

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

The optimal aim portfolio $aim(x_t) = Q^{-1}(q_0 + q^\top x_t)$ *of the GP investor can we written as a Markowitz portfolio where we replace the instantaneous expected stock return* $\mu_S(x_t) = \frac{1}{dt}\text{E}_t[dS_t] =$

$\mu_0 + \mu x_t$ *by the trading speed discounted value of the future stock expected returns:*

$$aim(x_t) = (\gamma\Sigma)^{-1} \int_0^\infty \omega_u \mathrm{E}_t[\mu_S(x_{t+u})]du \tag{49}$$

$$\omega_u = (\rho + \tau^\top)e^{-(\rho+\tau^\top)u} \tag{50}$$

Proof. The proof is in the appendix K ■

Note that by definition $\int_0^\infty \omega_u du = 1$, therefore we have that if $\kappa = 0$ then the optimal aim portfolio is the Markowitz portfolio. Only if there is some persistence in the factors that predict returns, is it optimal to deviate from the Markowitz portfolio. Of course, in the case where $\sigma_{xs} \neq 0$ then this result will no longer hold, as the investor will want to aim towards a portfolio that is also driven by its desire to hedge against variations in the investment opportunity set. The next section explores quantitatively the importance of these hedging demands.

In the general case it is possible to express the aim portfolio as follows:

$$aim(x_t) = Q^{-1}(q_0 + q^\top x_t) \tag{51}$$

$$= (\gamma\Sigma + q^\top \Omega q + 2\gamma\Sigma_{sx}q)^{-1} \int_0^\infty \omega_u \left\{ \mu_0 + \mu e^{-\kappa u}x_t - (\gamma\Sigma_{sx} + q^\top\Omega)(c_1 + c_2 e^{-\kappa u}x_t) \right\} du \tag{52}$$

$$\omega_u = (\rho + \tau^\top)e^{-(\rho+\tau^\top)u} \tag{53}$$

This allows us to interpret the hedging demands in three scenarios. First, if $\sigma_{xs} = 0$ and $\gamma_x = 0$, then $\Omega = 0$, and we recover the CMV preferences. Second, if $\sigma_{xs} = 0$ and $\gamma_x \neq 0$, then in the absence of transaction costs, it is optimal to hold the CMVE Markowitz portfolio (i.e., there are no hedging demands). However, with transaction costs, we do deviate from both the Markowitz portfolio and the CMV aim portfolio. Finally, if $\sigma_{xs} \neq 0$ and there are no transaction costs, then it is optimal to deviate from the Markowitz portfolio because of hedging demands. The optimal portfolio becomes $(\gamma\Sigma)^{-1}(\mu_0 + \mu x_t) - \Sigma^{-1}\Sigma_{sx}(c_1 + c_2 x)$. In the presence of transaction costs, we have the same structure in the optimal portfolio, but we need to take the weighted average of the expected return of a similar portfolio. The equation is more difficult to interpret, especially in the multi-asset case (note also that the numerical values for the c_1, c_2 matrices are different with and without t-costs). Therefore we turn to some specific examples and numerical simulations to illustrate the predictions of the model.

4 Hedging Demand and Transaction Costs: Numerical Example

4.1 The one Asset and one predictor case

To illustrate the model's predictions we first focus on the one asset-one factor case (that is $N = K = 1$) for the case where $\mu_0 = 0$, that is there is one single stock S_t and one single predictor variable x_t with dynamics:

$$dS(t) = \mu\, x_t\, dt + \sigma_1 dZ_1(t) \tag{54}$$

$$dx_t = -\kappa\, x(t)dt + \sigma_{x1} dZ_1(t) + \sigma_{x2} dZ_2(t) \tag{55}$$

where $Z_i(t)$ are independent Brownian motion.

We can solve the optimal portfolio of the non-myopic agent in the stationary case using Theorem 8.

For the case where there are no transaction costs, that is when $\Lambda = 0$, we find the optimal portfolio can be decomposed into the CMVE portfolio and a hedging portfolio HP that is:

$$n_t = CMVE_t + HP_t \tag{56}$$

$$CMVE_t = \frac{\mu}{\gamma \sigma_1^2}\, x_t \tag{57}$$

$$HP_t = -\frac{2(\frac{\mu}{\sigma_1})^2 \frac{\sigma_{x1}}{\sigma_1}}{\gamma(2\kappa + \rho + 2\frac{\mu}{\sigma_1}\sigma_{x1} + \sqrt{(2\kappa + \rho + 2\frac{\mu}{\sigma_1}\sigma_{x1})^2 + 4\frac{\gamma_x}{\gamma}(\frac{\mu}{\sigma_1})^2 \sigma_{x2}^2}}\, x_t \tag{58}$$

As expected the non-myopic agent deviates from the CMVE portfolio if and only if $\sigma_{x1} \neq 0$. She invests more in the stock the more negative is the covariance between x_t and S_t.

Turning now to the case with t-costs (setting $\Lambda_{11} = \lambda^2 > 0$), we first focus on the agent with CMV preferences. Recall that absent t-costs this agent would have myopic (locally mean-variance preferences), but deviates from that portfolio in the presence of t-costs.

Applying Theorem 9 we can derive the optimal aim portfolio and trading speed as follows:

$$aim_t^{CMV} = \frac{\rho + \tau}{\rho + \tau + \kappa} \frac{\mu}{\gamma \sigma_1^2} \tag{59}$$

$$\tau = \frac{1}{2}\left(\sqrt{\rho^2 + 4\gamma \frac{\sigma_1^2}{\lambda^2}} - \rho\right) \tag{60}$$

We note that the CMV-agent's aim portfolio always holds the CMVE portfolio only if $\kappa = 0$, otherwise her holdings are strictily decreasing in κ and increasing in $\rho + \tau$. The trading speed $\tau \in (0, \infty)$ is strictly increasing in $\gamma(\frac{\sigma_1}{\lambda})^2$.

Note that her optimal trading strategy, that is both the aim portfolio and trading speed, are independent of the covariance matrix of x_t, in that the CMV-agent would trade identically if x_t were deterministic (that is if $\sigma_{x1} = \sigma_{x2} = 0$).

Instead, if we consider the non-myopic agent with CARA with respect to both return and expected return shocks, applying theorem 8 we find that her aim portfolio and trading speeds are given by:

$$aim_t = \frac{\rho+\tau}{\rho+\tau+\kappa}\frac{\mu-(\gamma\sigma_1\sigma_{x1}+\Omega q)c_2}{\gamma\sigma_1^2+q^2\Omega+2\gamma\sigma_1\sigma_{x1}q} \tag{61}$$

$$\tau = \frac{1}{2}\left(\sqrt{\rho^2+4\left\{\gamma(\frac{\sigma_1}{\lambda}+\frac{q}{\lambda}\sigma_{x1})^2+\frac{q^2}{\lambda^2}\gamma_x\sigma_{x2}^2\right\}}-\rho\right) \tag{62}$$

$$c_2 = \frac{\sqrt{(2\kappa+\rho)^2+4\Omega\frac{q^2}{\lambda^2}}-\rho-2\kappa}{2\Omega} \tag{63}$$

where q is the constant that solves the following non-linear equation:[14]

$$c_2(\gamma\sigma_1\sigma_{x1}+\Omega q)+q(\rho+\kappa+\tau)=\mu \tag{64}$$

We see that, unlike for the CMV-agent, the non-myopic agent's optimal aim portfolio and trading speed are affected by the covariance matrix of the expected returns. In particular, her aim portfolio may actually hold more stock than the CMVE portfolio. We illustrate that with a few figures.

In figures 1-3 we compare trading strategies corresponding to different objective functions and for different sets of parameters. We are particularly interested in how the hedging demand of a non-myopic investor shapes her optimal trading strategy in the presence of transaction costs. Thus we report the trading strategy of an CMV investor who has the objective function (used by GP and CDS among others) given in (34), which is known to be myopic in the absence of transaction costs, and compare it with that of a source dependent risk-aversion investor (CARA) who maximizes (32) with $\gamma_x = \gamma$ and thus is risk-averse with respect to changes in the investment opportunity set.[15]

Figure 1 reports results for low trading costs and for a positive expected return signal ($x_0 = 1$). As expected, it shows that the CMV investor's optimal aim portfolio is very close to the mean-variance efficient Markowitz portfolio. Further, the CMV investor's strategy is independent of the correlation coefficient between the expected return signal and price changes. Instead, we see that for low transaction costs the CARA investor chooses a portfolio very similar to that of the classic no-transaction-cost *Merton* solution. Specifically, she displays a very large and positive hedging demand for the asset when correlation between x and dS becomes negative. This because the investor invests for the long run and perceives stock returns to be less risky for the long-run due to the negative correlation between expected returns and stock price changes. With negative correlation, expected returns changes offer a natural hedge for shocks to stock prices.

When we increase trading costs to more realistic levels, we see in figure 2 that the CMV investor chooses an aim portfolio that is uniformly lower than the Markowitz portfolio across all the correlation range. Intuitively, because of transaction costs the investor has to trade slowly into

[14]Note that the equation admits a strictly positive solution for any $\mu > 0$, since the left-hand side equals zero when $q = 0$ and tends to infinity when $q \to \infty$ (see Appendix J for details).

[15]Note that since we assume $\gamma_\mathcal{T} = 0$, the investor we consider is risk-neutral with respect to horizon realization risk.

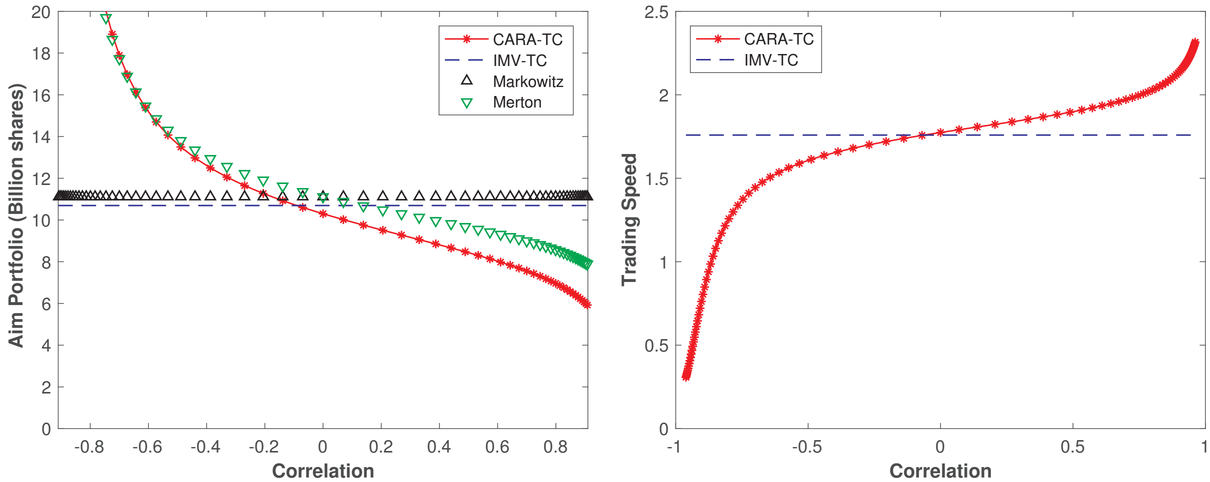

Figure 1: Parameters: $\mu_0 = 0$, $\mu = 1$, $\kappa = 0.1$, $\sigma_s = 0.3$, $\sigma_x = 0.1$, $\Lambda = 2 \times 10^{-11}$, $\gamma_x = 10^{-9}$, $\gamma = 10^{-9}$, $x_0 = 1$, $\rho = 0.8$.

her desired stock position. Because the signal also decays at rate $\kappa > 0$, it follows from theorem 9, that it is optimal to aim for a smaller position as the effective expected return that will be earned over the 'average' horizon of the position is lower than in the absence of transaction costs or with more persistent expected returns. This insight, which was also at the heart of GP's original paper, carries over for the non-myopic CARA investor, but only for positive correlation coefficients. Instead figure 2 shows that, surprisingly, when the correlation between signal and price change is sufficiently negative, the hedging demand can actually lead the investor to want to aim for a **larger** position in the risky asset than she would have chosen in the absence of transaction costs. We see on the picture that the point where the CARA-TC aim portfolio is larger than the Merton solution occurs for a correlation coefficent (between dS and dx) around -60%. Panel two on the same figure also shows that this coincides with a very steep drop in the trading speed. Instead, the CMV investor chooses the same constant trading speed irrespective of the level of the correlation coefficient.

Our results suggest that if the correlation between stock returns and their expected growth rates is sufficiently negative, then a long-term investor will want to hold more risky stocks in the presence of transaction costs than without, even though the expected return is decaying over time. At the time the investor will want to trade at a much lower speed than if she were myopic.

Our intuition for this surpising result is that, because of the negative correlation, the investor expects a lower expected return following a positive shock to stock prices and thus wants to trade out of stocks. Conversely, she will want to trade into stocks following a negative price shock. The aim portfolio is set so as to optimally trade-off the utility cost of deviating from the first-best portfolio and the transaction costs. When evaluating the cost of additional trading, the long-term agent weights these with her marginal utility. Thus costs that paid following the negative stock price shock will be weighted more. Therefore it can be optimal to aim for a higher stock position

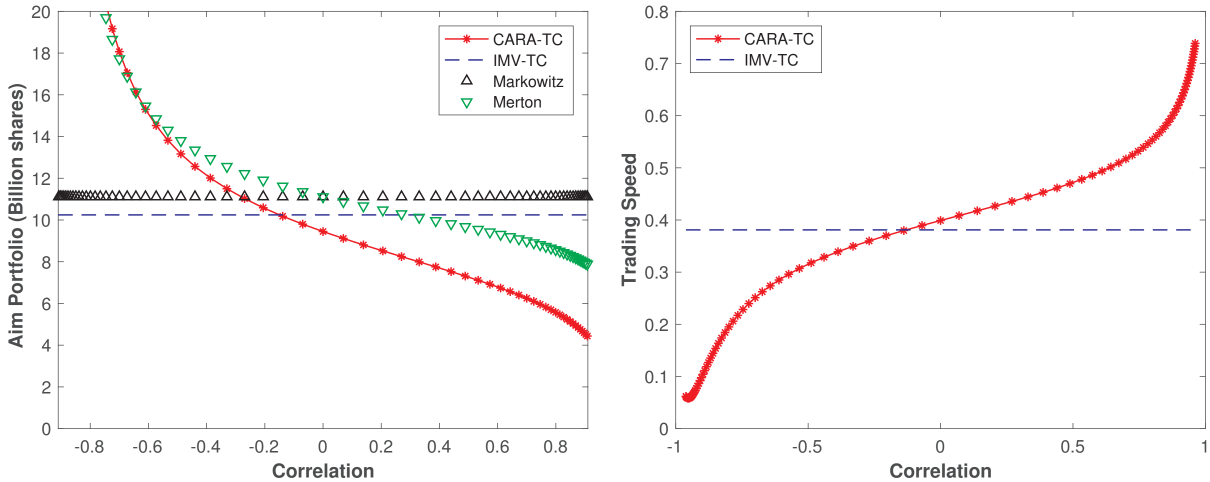

Figure 2: Parameters: $\mu_0 = 0$, $\mu = 1$, $\kappa = 0.1$, $\sigma_s = 0.3$, $\sigma_x = 0.1$, $\Lambda = 2 \times 10^{-10}$, $\gamma_x = 10^{-9}$, $\gamma = 10^{-9}$, $x_0 = 1$, $\rho = 0.8$.

and trade less to avoid paying the transaction costs in the high marginal utility states (following a negative stock price shock).

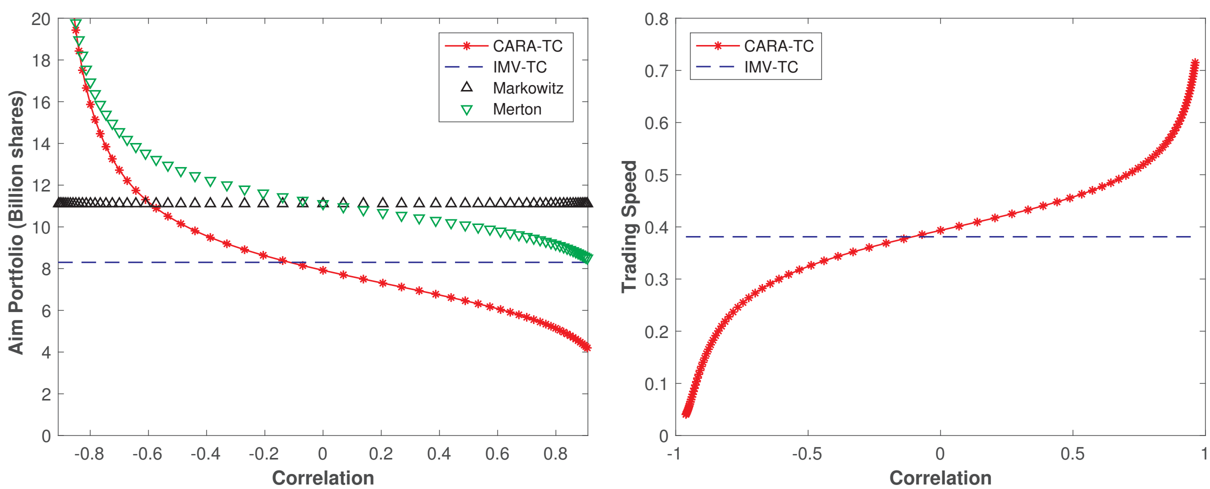

Figure 3: Parameters: $\mu_0 = 0$, $\mu = 1$, $\kappa = 0.4$, $\sigma_s = 0.3$, $\sigma_x = 0.1$, $\Lambda = 2 \times 10^{-10}$, $\gamma_x = 10^{-9}$, $\gamma = 10^{-9}$, $x_0 = 1$, $\rho = 0.8$.

In Figure 3 we show the effect of having a less persistent signal (with a higher κ). Since the expected return decays faster, the effective expected return earned over the life of the position decreases. Thus the myopic CMV investor decreases her position more relative to the Markowitz portfolio, which is unchanged. Similarly, the CARA investor in the presence of transaction costs reduces her position relative to the no t-cost Merton solution. Still we see that because of hedging demands, for sufficiently negative correlation (close to -80% in this case), the CARA investor's aim portfolio becomes larger than what she would choose in the absence of transaction costs. So even for fast decaying parameters, the hedging demands affect the optimal position of the long-term investor significantly.

The hedging demand of a non-myopic investor leads to a significantly different trading strategy than for a myopic investor in the presence of transaction costs. Below, we quantify with a realistic calibration the utility-based cost for a long-term investor of not properly accounting for the hedging demand in the presence of transaction costs.

4.2 The two asset and one predictor case

To illustrate the role of hedging demands in shaping the optimal portfolio choice we consider a very specific setup with two stocks with dynamics given by

$$dS_1(t) = \mu\, x_t\, dt + \sigma_1 dZ_1(t) \tag{65}$$

$$dS_2(t) = \sigma_2 dZ_2(t) \tag{66}$$

$$dx_t = -\kappa\, x(t)dt + \sigma_{x2} dZ_2(t) \tag{67}$$

where $Z_i(t)$ are independent Brownian motion. We further assume that the transaction cost matrix is diagonal with $\lambda_{11} = \lambda_1^2$ and $\Lambda_{22} = \lambda_2^2$. This is a special case of our general framework. We can solve the optimal portfolio of the non-myopic agent in the stationary case using Theorem 8. For the case where there are no transaction costs, that is when $\Lambda = 0$, we find the optimal portfolio can be decomposed into the CMVE portfolio that only loads on asset 1, and a hedging portfolio HP given by:

$$n_t = CMVE_t + HP_t \tag{68}$$

$$CMVE_t = [\frac{\mu\, x_t}{\gamma \sigma_1^2}; 0]^\top \tag{69}$$

$$HP_t = [0; -\frac{\mu^2 \sigma_{x2}}{\gamma \sigma_1^2 \sigma_2 (\rho + 2\kappa)} x_t]^\top \tag{70}$$

We see that the myopic agent only trades asset 1, but has no demand for asset 2, since the latter has zero expected (excess) return and positive variance. It therefore does not improve the conditional mean-variance efficient frontier. Instead, asset 2 is correlated with asset 1's expected return and therefore is desirable to trade for a non-myopic agent as it allows to hedge against changes in the investment opportunity set. Indeed, we see that the hedging portfolio goes long asset 2 if it is negatively correlated with asset 1's expected return and shorts it otherwise. Since asset 1 has zero correlation with its expected return, it is not traded for hedging purposes.

Our example is engineered such that each asset is uniquely associated with the CMVE and the HP portfolios respectively. We now turn to the case with transaction costs (with $\lambda_i > 0\ \forall i = 1, 2$) to see how the assets enter the aim portfolio.

We start with the CMV-utility agent. Applying Theorem 9 we can derive the optimal aim portfolio and trading speed as follows:

$$aim_t^{CMV} = [\frac{\rho+\tau_{11}}{\rho+\tau_{11}+\kappa}\frac{\mu}{\gamma\sigma_1^2};\ 0]^\top \tag{71}$$

$$\tau_{11} = \frac{1}{2}(\sqrt{\rho^2+4\gamma\frac{\sigma_1^2}{\lambda_1^2}}-\rho) \tag{72}$$

$$\tau_{22} = \frac{1}{2}(\sqrt{\rho^2+4\gamma\frac{\sigma_2^2}{\lambda_2^2}}-\rho) \tag{73}$$

$$\tau_{12} = \tau_{21} = 0 \tag{74}$$

Since the aim portfolio for a CMV-investor is the trading-speed discounted value of the future expected CMVE portfolios and given that the latter only hold asset 1, we see that the aim portfolio only comprises asset 1 as well. The trading speed is a diagonal matrix, which implies that positions in asset 2 do not affet how to optimally trade asset 1. Instead, the optimal strategy is for the agent to trade out of any initial position she might have in asset 2 at a constant trading speed and towards 0, the optimal position for asset 2 in the CMVE portfolio. Thus for a myopic-CMV agent, trading in asset 2 occurs only in as much as she would be endowed with a non-zero position in that asset. As in the case without trading costs, there is no motive for trading (or holding) asset 2 in the case with transaction costs. We also see, consistent with our general results, that the CMV-agent's optimal aim and trading speed are not affected by the covariance matrix of the expected return variable x_t.

We now turn to the optimal aim portfolio for non-myopic CARA agent, and show that for such an agent with a long-horizon, the variance of expected returns and its covariance with the underlying stocks dramatically affect her optimal holdings in the aim portfolio as well as the trading speed matrix.

Solving the system for the optimal aim portfolio and trading speed reduces to a system of non-linear equation:

$$aim_t = Q^{-1}q \tag{75}$$

$$Q = \Lambda\tau \tag{76}$$

$$\tau_{11} = \frac{1}{2}(\sqrt{\rho^2+4\gamma\frac{\sigma_1^2}{\lambda_1^2}+4\gamma\frac{q_1^2}{\lambda_1^2}\sigma_{x2}^2-4\frac{\lambda_1^2}{\lambda_2^2}\tau_{12}^2}-\rho) \tag{77}$$

$$\tau_{22} = \frac{1}{2}(\sqrt{\rho^2+4\gamma(\frac{\sigma_2}{\lambda_2}+\frac{q_2}{\lambda_2}\sigma_{x2})^2-4\frac{\lambda_1^2}{\lambda_2^2}\tau_{12}^2}-\rho) \tag{78}$$

$$\tau_{12} = \frac{\lambda_2^2}{\lambda_1^2}\tau_{21} = \frac{\lambda_2}{\lambda_1}\frac{\gamma\frac{\sigma_2 q_1}{\lambda_2\lambda_1}\sigma_{x2}+\gamma\frac{q_1 q_2}{\lambda_1\lambda_2}\sigma_{x2}^2}{\rho+\tau_{11}+\tau_{22}} \tag{79}$$

and where the q_1, q_2, c_2 solve

$$c_2 = \frac{-\rho - 2\kappa + \sqrt{(\rho+2\kappa)^2 + 4(\frac{q_1^2}{\lambda_1^2} + \frac{q_2^2}{\lambda_2^2})\gamma\sigma_{x2}^2}}{2\gamma\sigma_{x2}^2} \tag{80}$$

$$\mu = q_1(\kappa + \rho + c_2\gamma\sigma_{x2}^2 + \tau_{11}) + q_2\frac{\lambda_1^2}{\lambda_2^2}\tau_{12} \tag{81}$$

$$-c_2\sigma_2\gamma\sigma_{x2} = q_2(\kappa + \rho + c_2\gamma\sigma_{x2}^2 + \tau_{22}) + q_1\tau_{12} \tag{82}$$

We solve this system numerically and show how aim portfolio changes with parameter of the model, and in particular with the diffusion coefficients of x_t. We specifically compare the optimal solution for the non-myopic CARA agent with the benchmarks we have specified earlier. Figure 4 illustrates the aim portfolios in asset 1 (left panel) and asset 2 (right panel) as σ_{x2} is varied. For asset 1, the aim portfolios are the same for CMV (CMVE) and non-myopic CARA (Merton) agent. For asset 2, both CMV and CMVE agents have zero desired position whereas non-myopic CARA and the Merton agents differ considerably due to the positive or negative σ_{x2}. We find that relative to Merton solution, non-myopic CARA agent shrinks her portfolio slightly to zero to mitigate the impact of transaction costs. Figure 5 illustrates the aim portfolios in asset 1 (left panel) and asset 2 (right panel) as κ is varied while keeping $\sigma_{x2} = -0.5$. For asset 1, the aim portfolios are the same and constant for CMVE and Merton agents while CMV and non-myopic CARA reduce their position in asset 1 as κ increases. Further, while the Merton agent also reduces her position in asset 2, non-myopic CARA agent responds more dramatically as κ is increased.

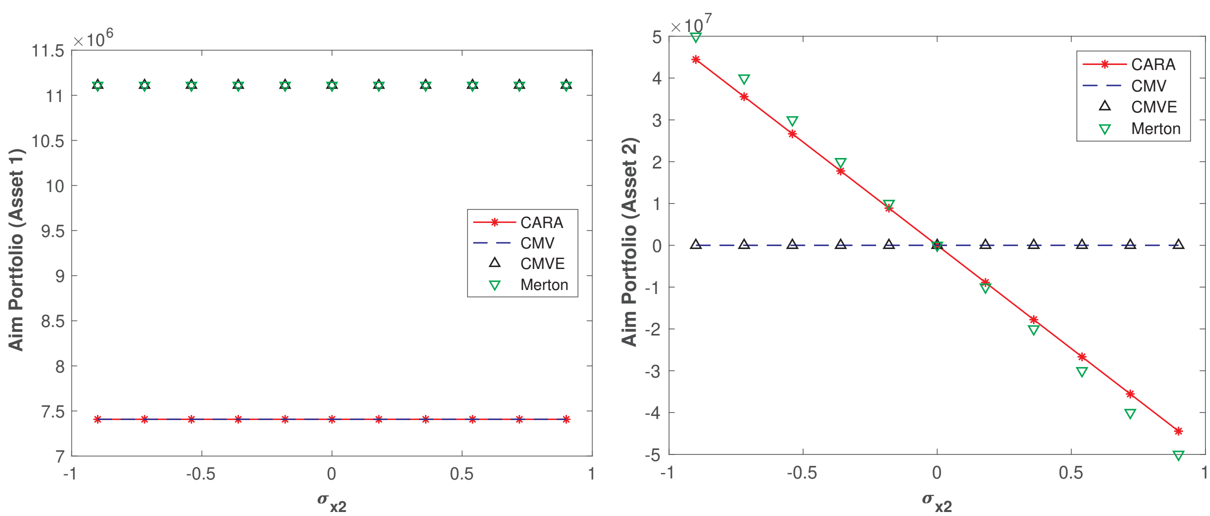

Figure 4: Parameters: $\mu_0 = 0$, $\mu = 1$, $\kappa = 0.5$, $\sigma_1 = 0.3$, $\sigma_2 = 0.1$, $\sigma_x = 0.2$, $\Lambda = 0.01\Sigma$, $\gamma_x = 10^{-6}$, $\gamma = 10^{-6}$, $x = 1$, $\rho = 1$.

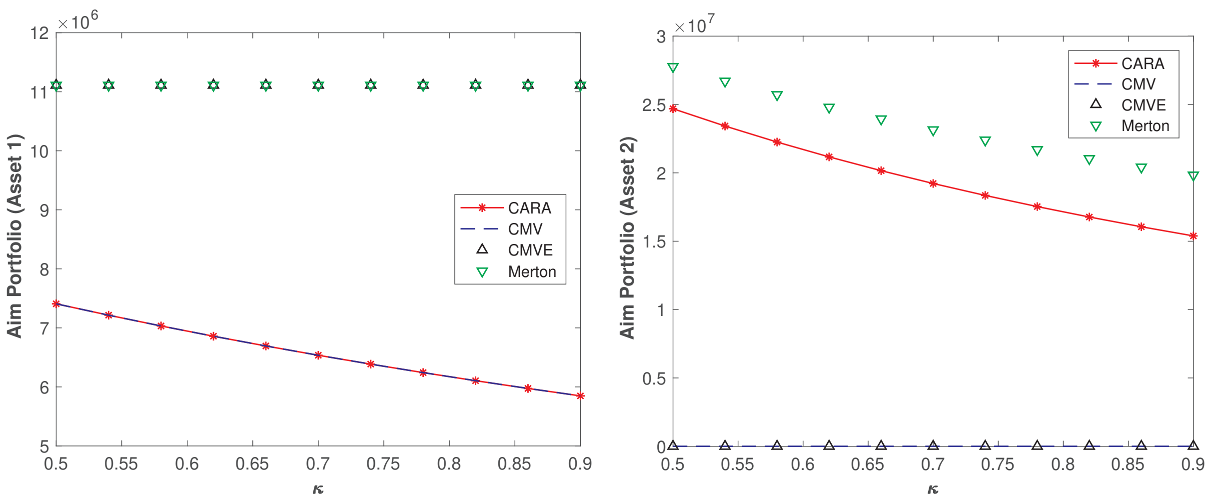

Figure 5: Parameters: $\mu_0 = 0$, $\mu = 1$, $\kappa = 0.5$, $\sigma_1 = 0.3$, $\sigma_2 = 0.1$, $\sigma_x = 0.2$, $\sigma_{x2} = -0.5$, $\Lambda = 0.01\Sigma$, $\gamma_x = 10^{-6}$, $\gamma = 10^{-6}$, $x = 1$, $\rho = 1$.

5 Empirical Application with the Informativeness of Retail Order Flow

In this section, we implement our methodology based on the recent findings that net retail order flow predicts future returns at the stock level. (Boehmer, Jones, Zhang, and Zhang, 2021) propose an easy algorithm to identify marketable retail purchases and sales and find that individual stocks with net buying by retail investors outperform stocks with net sells over the following week. We will use retail order imbalance at the stock level as a predictor for next day returns.

We realistically calibrate transaction costs based on a large institutional order data. We illustrate that there are economically significant benefits of using our approach for a CARA investor in an out-of-sample experiment compared to a CMV investor who ignores the correlation between the innovations in stock prices and predictors.

5.1 Model calibration

Using three-year data between 2014-01-01 and 2016-12-31, we first test the relation between net retail order flow and subsequent daily returns. We use two stocks from the top 5 stocks having the largest market capitalization as of 2013-12-31, Johnson and Johnson (JNJ) and Exxon Mobil (XOM). We will assume that our first asset is JNJ and second asset is XOM in our matrix notations going forward. We will first calibrate the model using percentage returns and then scale them appropriately in our out-of-sample experiment. We will use the superscript scl to denote that this parameter will be scaled.

Let $N^b_{i,t}$ ($N^s_{i,t}$) be the number of retail buy (sell) orders on stock i and day t, using the the retail trade classification of Boehmer, Jones, Zhang, and Zhang (2021). Our return predictor at the stock level is then given by

$$x_{i,t} = \frac{N^b_{i,t} - N^s_{i,t}}{N^b_{i,t} + N^s_{i,t}}.$$

We then run the following regressions with our two stocks to test the empirical relation between net retail order flow and subsequent daily returns:

$$r_{1,t+1} = \beta_1 x_{1,t} + \epsilon_{1,t+1} \tag{83}$$

$$r_{2,t+1} = \beta_2 x_{2,t} + \epsilon_{2,t+1} \tag{84}$$

$$\tag{85}$$

Table 1 reports the regression results where standard errors are adjusted for heteroskedasticity. In both specifications, the coefficients are quite similar in magnitude with being close to 60 bps and they are both significant at 10% level. Using these regression results, we calibrate μ^{scl} to be $\begin{bmatrix} 0.0058 & 0 \\ 0 & 0.0057 \end{bmatrix}$. Since the constants in Table 1 are insignificant, we will set μ_0 to be the zero vector, i.e., $\begin{bmatrix} 0 \\ 0 \end{bmatrix}$.

	Dependent variable:	
	$r_{1,t+1}$	$r_{2,t+1}$
$x_{1,t}$ (JNJ)	0.0058*	
	(0.0032)	
$x_{2,t}$ (XOM)		0.0057*
		(0.0034)
Constant	0.0002	0.00005
	(0.0004)	(0.0004)
Observations	755	755
Adjusted R^2	0.0030	0.0025

Table 1: Net retail order flow imbalance and subsequent returns

We then compute the variance-covariance matrix of the daily returns, $\Sigma^{\mathsf{scl}} = \text{Var}(\epsilon)$ where $\epsilon = [\epsilon_1 \ \epsilon_2]$. Using this estimation, we calibrate Σ^{scl} to be $10^{-4} \times \begin{bmatrix} 0.874 & 0.537 \\ 0.537 & 1.513 \end{bmatrix}$.

We estimate the mean-reversion parameters of the predictors by running the following regressions at the stock level:

$$\Delta x_{1,t+1} = -\kappa_1 x_{1,t} + \varepsilon_{1,t+1} \tag{86}$$

$$\Delta x_{2,t+1} = -\kappa_2 x_{2,t} + \varepsilon_{2,t+1} \tag{87}$$

Table 2 reports the regression results. For both stocks, the net retail order flow exhibits strong mean-reversion characteristics with statistically significant $AR(1)$ coefficients. Our results show that this stock-level signal is quite fast-decaying. For JNJ stock, the half-life of the net retail order imbalance signal is approximately one-day while for XOM stock the half-life of the signal is 1.6 days. This analysis leads to the calibration of κ with $\begin{bmatrix} 0.4862 & 0 \\ 0 & 0.3469 \end{bmatrix}$.

Σ_x can be estimated by computing the variance-covariance matrix of the innovations in the predictors, ε_{t+1}. We obtain Σ_x to be $\begin{bmatrix} 0.0108 & 0.0048 \\ 0.0048 & 0.0123 \end{bmatrix}$. Finally, $\Sigma_{sx}^{\mathsf{scl}}$ is given by the sample covariance between ϵ_{t+1} and ε_{t+1} which yields the following matrix:

$$\Sigma_{sx}^{\mathsf{scl}} = 10^{-4} \times \begin{bmatrix} -3.425 & -2.042 \\ -3.071 & -5.897 \end{bmatrix}.$$

The negative values in this covariance matrix are striking. At the individual stock level, it suggests

	Dependent variable:	
	$\Delta x_{1,t+1}$	$\Delta x_{2,t+1}$
$x_{1,t}$ (JNJ)	-0.4862^{***}	
	(0.031)	
$x_{2,t}$ (XOM)		-0.3469^{***}
		(0.028)
Constant	0.020^{***}	0.002
	(0.004)	(0.004)
Observations	755	755
Adjusted R^2	0.243	0.173

Table 2: AR(1) regressions of net retail order flow imbalance

that when there is a positive shock to the price of either JNJ or XOM, there will be less retail buying activity on these stocks contemporaneously. This is consistent with a contrarian trading strategy at the aggregate retail level, i.e., retail traders tend to sell (buy) a stock with positive (negative) daily returns.

5.2 Calibration of the Transaction Costs

To calibrate the transaction cost multipliers of our model realistically, we use proprietary execution data from the historical order databases of a large investment bank. The orders primarily originate from institutional money managers who would like to minimize the costs of executing large amounts of stock trading through algorithmic trading services. The data consists of two frequently used trading algorithms, volume weighted average price (VWAP) and percentage of volume (PoV). The VWAP strategy aims to achieve an average execution price that is as close as possible to the volume weighted average price over the execution horizon. The main objective of the PoV strategy is to have constant participation rate in the market along the trading period.

The 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%. We use the top 50 stocks in terms of market capitalization to estimate the price impact for a liquid subset.

The standard measure of institutional trading costs is given by implementation shortfall (IS). It is computed as the normalized difference between the average execution price and the mid-quote price of the asset prior to the start of the execution. Formally, the IS of the ith parent-order is given by

$$IS_i = D_i \frac{P_i^{\text{avg}} - P_{i,0}}{P_{i,0}}, \tag{88}$$

where Q_i is the order size (in shares) with $Q_i > 0$ ($Q_i < 0$) for buy (sell) orders, P_i^{avg} is the volume-weighted execution price of the parent-order, D_i equals 1 (-1) if the order is buy (sell), and $P_{i,0}$ is the mid-quote price of the security (arrival price) when the parent order starts being executed.

We first estimate the price impact coefficient as a function of the ratio of the order size to the daily volume using the complete data set. Formally, we run the following regression with our order data from 50 largest stocks:

$$IS_i = \theta \frac{Q_i}{DayVlm_i} + \varepsilon_i$$

where Q_i is the number of shares to be bought or sold, $DayVlm_i$ is the daily volume of the stock. Here, $\frac{Q_i}{DayVlm_i}$ measures the size of the order as a fraction of the total daily volume. We measure IS in basis points.

Table 3 illustrates the estimated coefficient, θ. The reported standard errors are clustered at the calendar day level. We find that θ is statistically significant with a t-statistic of 3.9. The economic magnitude is also large. For an order that aims to trade 1% of daily volume, the expected transaction cost is 3.6 bps.

	Dependent variable:
	IS (bps)
$\frac{Q}{DayVlm}$	363.20***
	(91.66)
Constant	0.52
	(1.12)
Observations	16,532
Adjusted R^2	0.001

Table 3: Estimation of the price impact from institutional order data set

According to our quadratic transaction cost model, trading Q shares of stock i would move its (average) price by $\frac{\lambda_i Q}{2}$ where λ_i is its price impact coefficient. Since the typical daily volume of the stocks will differ, we can calibrate the price impact coefficients at the stock level by utilizing the average daily volume of each stock in the sample period. Let $\hat{V}_{i,t}$ be the daily volume for stock i on day t. Then, $\lambda_{i,t}^{\mathsf{scl}}$ would be given by $\frac{2\theta}{V_{i,t}}$ on a single calendar day t. We set λ_i^{scl} to be the sample average across all trading days, i.e., $\lambda_i^{\mathsf{scl}} = \frac{2\theta}{T} \sum_{t=1}^{T} \frac{1}{V_{i,t}}$. Put differently, $\lambda_i^{\mathsf{scl}} = 2\theta \hat{Z}_i$ where $\hat{Z}_i$ is a measure of average turnover given by $\frac{1}{T}\sum_{t=1}^{T}\frac{1}{V_{i,t}}$. For JNJ (XOM), we have $\hat{Z}_1 = 1.43 \times 10^{-7}$ ($\hat{Z}_2 = 8.8 \times 10^{-8}$) suggesting that JNJ's price impact coefficient is roughly 60% larger than that of XOM. Computing the λ for JNJ and XOM, we obtain the calibration of Λ^{scl} in matrix notation:

$$\Lambda^{\mathsf{scl}} = \begin{bmatrix} \lambda_1^{\mathsf{scl}} & 0 \\ 0 & \lambda_2^{\mathsf{scl}} \end{bmatrix} = 10^{-3}\begin{bmatrix} 0.1040 & 0 \\ 0 & 0.0645 \end{bmatrix}.$$

5.3 Out-of-sample experiments

In this section, we compare the trading policies for CARA and CMV agents for several non-overlapping out-of-sample investment horizons. As we noted in the previous section, our model parameters will be scaled appropriately with the initial stock price observed at the beginning of each investment horizon, j.[16] Let the initial stock prices equal $S_0^j = \begin{bmatrix} S_{1,0}^j \\ S_{2,0}^j \end{bmatrix}$ for this out-of-sample period j. Both agents will then use the following scaled parameters in the jth out-of-sample period:

$$\mu^j = \mu^{\mathsf{scl}}\,\mathrm{diag}(S_0^j) \tag{89}$$

$$\Sigma^j = \mathrm{diag}(S_0^j)\Sigma^{\mathsf{scl}}\,\mathrm{diag}(S_0^j) \tag{90}$$

will be observed and the agents will trade according

5.4 Aim Portfolios

Using these parameters, we first compute the aim portfolios corresponding to the traditional Markowitz and Merton portfolios which ignore trading costs and the aim portfolios corresponding to the CARA and the CMV investor in the presence of trading costs. Figure 6 illustrates the aim portfolios of each policy as a function of x_t. Note that since μ_0 is equal to zero, the aim portfolios cross at the origin. Given the negative correlation, we observe that Merton portfolio has higher positive slope compared to the Markowitz portfolio.

5.5 Performance of CARA and CMV Policies

We use the filtered values of x_t and divide the sample into 112 trading intervals of 126 trading days (i.e., six-months) to evaluate the performance of the CARA and CMV policies in the presence of trading costs. At the beginning of each interval, we start with $0 and trade according to trading policies determined by CARA and CMV policies, using the previously estimated return and t-cost parameters. We compute the stock position in shares and the total accumulated wealth after subtracting the trading costs paid every day. Utilizing all of the samples, we can compute the average utility corresponding to each policy and the resulting certainty equivalent wealth for both policies. Below we estimate the utility cost for a CARA investor of following the 'sub-optimal' policy of an CMV investor (who effectively ignores Merton-style hedging demand).

[16]For $j > 1$, S_0^j will be the last price observed in the previous period, i.e., S_T^{j-1}.

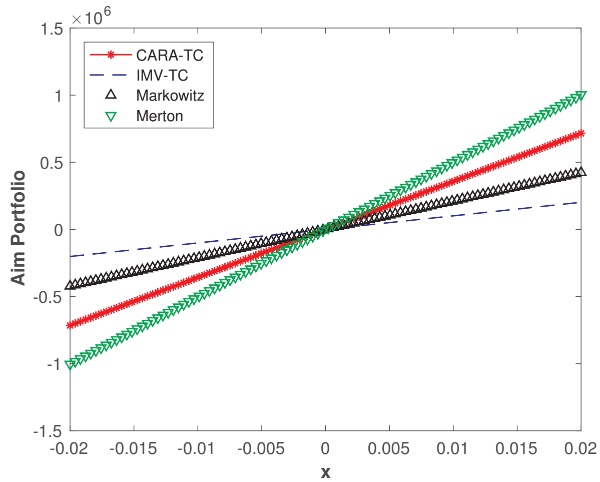

Figure 6: Aim portfolios of four different policy as a function of x_t.

Table 4: Policy comparison.

Statistic	CARA-TC	CMV-TC	Diff
Avg Util	-0.9993	-0.9995	1.96×10^{-4}
S.E	1.3×10^{-4}	0.9×10^{-4}	3.8×10^{-5}
CE	72,228	52,569	19,659

Figure 7 compares the CARA policy to the CMV policy in a single path. The top-left panel illustrates that the difference in wealth is increasing as time elapses. The top-right panel suggests that the CARA policy is more aggressive compared to the CMV policy. This is driven by the relatively higher slope of the CARA policy that we observed in the comparison of the aim portfolios in the prior section. We can see that this aggressiveness leads to higher trading costs for the CARA investor but this loss is compensated with higher profits made from the aggressive position as illustrated in the top-left panel.

We find that these observations hold for the average of all paths as well. Figure 8 compares the CARA policy to the CMV policy by taking the average of all statistics across 112 trials. The top-left panel illustrates on average the wealth is higher for the CARA investor. The top-right panel suggests that the CARA policy is again more aggressive compared to the CMV policy on average basis. The bottom panel illustrates the cumulative certainty-equivalent wealth across all samples and we find that CARA policy achieves significantly higher certainty equivalent throughout the 6-month horizon. As we can see in table 4 the certainty equivalent achieved by making use of hedging demands is 37% higher for the CARA investor. So the utility loss of using the t-cost optimal policy of a myopic investor, that does not adjust the aim portfolio and trading speed to account for the negative correlation between expected returns stock returns, is very substantial for long-term investor with standard expected utility.

6 Conclusion

In the presence of time-varying expected returns, long-term investors with CARA utility who ignore trading costs deviate from the conditional mean-variance efficient portfolio to hedge against the negative impact of the time variation in expected returns on the marginal utility of the investor. In the recent literature, the dynamic trading policy based on conditional mean-variance preferences that incorporates transaction costs has been very popular. Surprisingly, this trading policy has no hedging component. We propose a set of preferences based on stochastic differential utility with source-dependent risk-aversion, which nest the widely used conditional mean-variance and CARA 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 non-myopic 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 a realistic calibration where we time the S&P 500 return based on its filtered latent predicted expected return, hedging demands significantly affect strategy performance.

Figure 7: This figure compares the in-sample performance of the CARA policy with the CMV policy in the presence of trading costs for a single path. Both strategies start from zero-wealth. CARA policy is the optimal policy corresponding to an investor with CARA preferences. CMV policy is the optimal policy corresponding to an investor with CMV-preferences. Top-left panel shows the cumulative wealth of each policy. Top-right panel shows the each strategy's share position in the S&P 500 index. Bottom-left panel displays the filtered x_t values for the single path. Bottom-right panel illustrates the cumulative trading costs paid by each policy due to the rebalancing of the strategies.

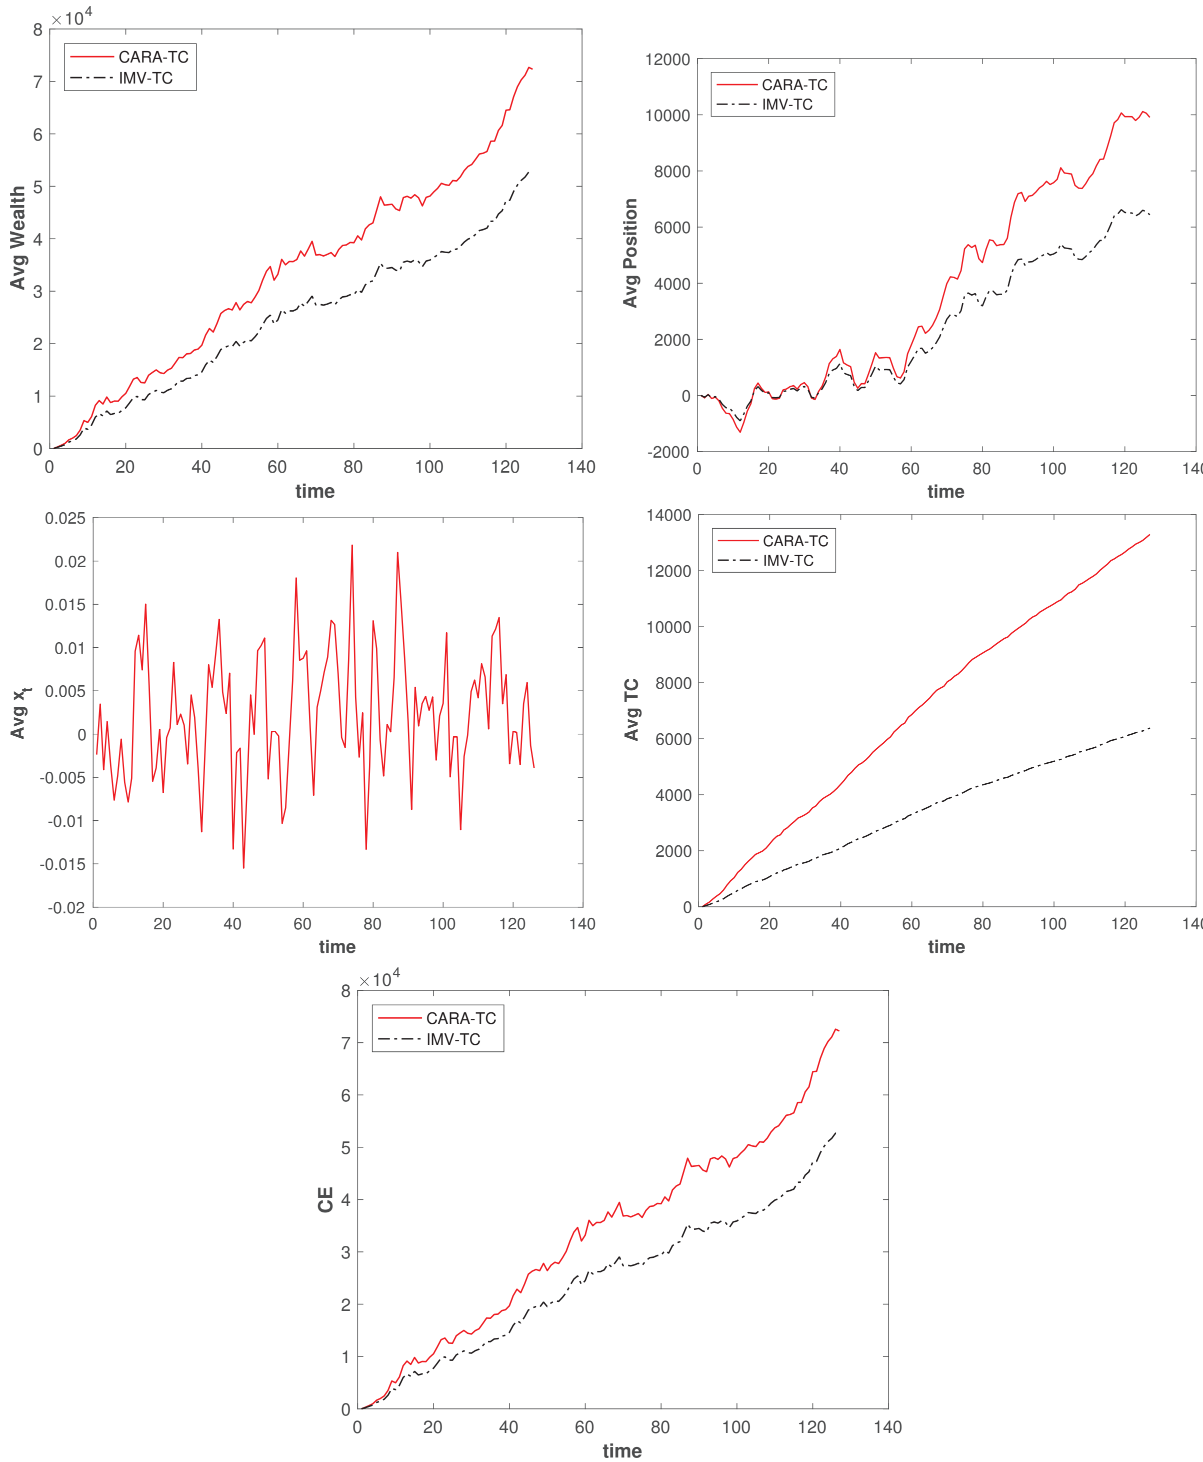

Figure 8: This figure compares the average performance of the CARA policy with the CMV policy in the presence of trading costs across all 112 6-month trading periods. Both strategies start from zero-wealth. CARA policy is the optimal policy corresponding to an investor with CARA preferences. CMV policy is the optimal policy corresponding to an investor with CMV preferences. Top-left panel shows the mean cumulative wealth of each policy. Top-right panel shows the each strategy's average share position in the S&P 500 index. Bottom-left panel displays the average filtered x_t values for the single path. Bottom-right panel illustrates the average cumulative trading costs paid by each policy due to the rebalancing of the strategies. The bottom panel displays the cumulative certainty-equivalent wealth.

References

Balduzzi, Pierluigi, and Anthony W Lynch, 1999, Transaction costs and predictability: Some utility cost calculations, *Journal of Financial Economics* 52, 47–78.

Boehmer, Ekkehart, Charles M Jones, Xiaoyan Zhang, and Xinran Zhang, 2021, Tracking retail investor activity, *The Journal of Finance* 76, 2249–2305.

Brandt, Michael W, 1999, Estimating portfolio and consumption choice: A conditional euler equations approach, *The Journal of Finance* 54, 1609–1645.

Brennan, Michael J., Eduardo S. Schwartz, and Ronald Lagnado, 1997, Strategic asset allocation, *Journal of Economic Dynamics and Control* 21, 1377–1403.

Brown, David B, and James E Smith, 2011, Dynamic portfolio optimization with transaction costs: Heuristics and dual bounds, *Management Science* 57, 1752–1770.

Campbell, John, 1999, Asset prices, consumption, and the business cycle, in John B. Taylor, and Michael Woodford, ed.: *Handbook of Macroeconomics* (North Holland: Amsterdam).

Campbell, John Y, and Luis M Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors* (Oxford University Press: Oxford).

Collin-Dufresne, Pierre, Kent Daniel, and Mehmet Sağlam, 2020, Liquidity regimes and optimal dynamic asset allocation, *Journal of Financial Economics* 136, 379–406.

Constantinides, George M., 1986, Capital market equilibrium with transaction costs, *The Journal of Political Economy* 94, 842–862.

Cox, John C, and Chi-fu Huang, 1989, Optimal consumption and portfolio policies when asset prices follow a diffusion process, *Journal of economic theory* 49, 33–83.

Danielsson, Jon, Hyun Song Shin, and Jean-Pierre Zigrand, 2012, Endogenous extreme events and the dual role of prices, *Annu. Rev. Econ.* 4, 111–129.

Davis, M. H. A., and A. R. Norman, 1990, Portfolio selection with transaction costs, *Mathematics of Operations Research* 15, 676–713.

Du, Songzi, and Haoxiang Zhu, 2017, What is the optimal trading frequency in financial markets?, *The Review of Economic Studies* 84, 1606–1651.

Duffie, Darrell, and Haoxiang Zhu, 2017, Size discovery, *The Review of Financial Studies* 30, 1095–1150.

Dumas, Bernard, and Elisa Luciano, 1991, An exact solution to a dynamic portfolio choice problem under transactions costs, *Journal of Finance* 46, 577–595.

Gârleanu, Nicolae, and Lasse H. Pedersen, 2013, Dynamic trading with predictable returns and transaction costs, *The Journal of Finance* 68, 2309–2340.

Gourinchas, Pierre-Olivier, Walker Ray, and Dimitri Vayanos, 2021, A preferred-habitat model of term premia, exchange rates and monetary policy spillovers, Unpublished working paper. UC Berkeley.

Greenwood, Robin, and Dimitri Vayanos, 2014, Bond supply and excess bond returns, *The Review of Financial Studies* 27, 663–713.

Hugonnier, Julien, Florian Pelgrin, and Pascal St-Amour, 2012, Health and (other) asset holdings, *The Review of Economic Studies* 0, 1–48.

Kim, Tong Suk, and Edward Omberg, 1996, Dynamic nonmyopic portfolio behavior, *Review of Financial Studies* 9, 141–161.

Litterman, Robert, 2005, Multi-period portfolio optimization, Unpublished working paper. Goldman Sachs.

Liu, Hong, 2004, Optimal consumption and investment with transaction costs and multiple risky assets, *The Journal of Finance* 59, 289–338.

Longstaff, Francis A., 2001, Optimal portfolio choice and the valuation of illiquid securities, *The Review of Financial Studies* 14, 407–431.

Lynch, Anthony W, and Pierluigi Balduzzi, 2000, Predictability and transaction costs: The impact on rebalancing rules and behavior, *The Journal of Finance* 55, 2285–2309.

Lynch, Anthony W, and Sinan Tan, 2010, Multiple risky assets, transaction costs, and return predictability: Allocation rules and implications for US investors, *Journal of Financial and Quantitative Analysis* 45, 1015–1053.

Malkhozov, Aytek, Philippe Mueller, Andrea Vedolin, and Gyuri Venter, 2016, Mortgage risk and the yield curve, *The Review of Financial Studies* 29, 1220–1253.

Markowitz, Harry M., 1952, Portfolio selection, *Journal of Finance* 7, 77–91.

Merton, Robert C, 1969, Lifetime portfolio selection under uncertainty: The continuous-time case, *The Review of Economics and Statistics* 51, 247–257.

Merton, Robert C., 1971, Optimum consumption and portfolio rules in a continuous-time model, *Journal of Economic Theory* 3, 373–413.

Shreve, Steven E, and H. Mete Soner, 1994, Optimal investment and consumption with transaction costs, *The Annals of Applied Probability* 4, 609–692.

Skiadas, Costis, 2008, Dynamic portfolio choice and risk-aversion, in J. R. Birge, and V. Linetsky, ed.: *Handbooks in Operations Research and Management Science*, vol. 15 . chap. 19, pp. 789–843 (Elsevier B.V.: Amsterdam).

Vayanos, Dimitri, and Jean-Luc Vila, 2021, A preferred-habitat model of the term structure of interest rates, *Econometrica* 89, 77–112.

APPENDIX

A Stochastic Differential Utility of Terminal Wealth

Consider an agent with a wealth process W_t who trades in a financial market, where the uncertainty is generated by a vector of independent Brownian motion $Z(t)$, and who has expected utility of terminal wealth with twice-differential, increasing and concave utility function $U(W_T)$. Note that by definition $M_t = \mathrm{E}_t[U(W_T)]$ is a martingale and therefore we may write:

$$dM_t = \sigma_M^\top dZ_t$$

Now define the certainty equivalent process $H_t = U^{-1}(M_t)$ which satisfies the boundary condition $H_T = W_T$. Defining

$$dH_t = \mu_H dt + \sigma_H^\top dZ_t \tag{91}$$

Then we have

$$dU(H_t) = \left(\frac{1}{2}U''(H)||\sigma_H||^2 + U'(H)\mu_H\right)dt + U'(H)\sigma_H^\top dZ_t$$

Since $M_t = U(H_t)$ comparing the two processes we get:

$$\mu_H = -\frac{1}{2}\frac{U''(H)}{U'(H)}||\sigma_H||^2 \tag{92}$$

It follows that we can define the certainty equivalent of an investor who has expected utility of terminal wealth as the solution (H_t, σ_H) of a backward-stochastic differential equation:

$$H_t = \mathrm{E}_t[W_T - \int_t^T \mu_H(H_t, \sigma_H)dt] \tag{93}$$

where the driver of the BSDE is given in equation (92) above.

To summarize, we have shown that, for an agent with an arbitrary wealth process W_t driven by a vector or N Brownian motions, who has expected utility of terminal wealth $\mathrm{E}[U(W_T)]$, we can define his certainty equivalent H_t in two different ways. First, the traditional definition $H_t = U^{-1}(\mathrm{E}_t[U(W_T)])$. Second, as the solution of the BSDE given in (92-93) above. Both are equivalent. It turns out the BSDE definition lends itself naturally to a generalization where the agent has source-dependent risk-aversion in that she attaches different risk-aversion to different Brownian motions.

B Source-Dependent SDU with Vanishing Risk Aversion to Expected return shocks

Specifically, consider the case of two vectors of independent Brownian motions Z^s, Z^x, then we define the certainty equivalent of our "source-dependent stochastic differential utility" agent who consumes only at maturity T, as the solution $(H_t, \sigma_{H,s}, \sigma_{H,x})$ of the following BSDE:

$$H_t = \mathrm{E}_t[W_T - \int_t^T \mu_H(H_t, \sigma_{H,s}, \sigma_{H,x})dt] \tag{94}$$

$$\mu_H = -\frac{1}{2}\frac{U_1''(H)}{U_1'(H)}||\sigma_{H,s}||^2 - \frac{1}{2}\frac{U_2''(H)}{U_2'(H)}||\sigma_{H,x}||^2 \tag{95}$$

where two different (twice-differential, strictly increasing and concave) utility functions U_i $i = 1, 2$ apply to the different sources of diffusion risk.[17] Of course, if we pick $U_1 = U_2$, then H_t is simply the standard certainty equivalent of an agent that has expected utility of terminal wealth as shown in the previous section. Otherwise, we define H_t as the certainty equivalent of an agent that has source-dependent risk-aversion and applies different risk-aversion to different sources of diffusion risk.

For CARA utility functions $U_i(w) = -e^{-\gamma_i w}$ $\forall i = 1, 2$, we obtain the following expression for the BSDE satisfied by the certainty equivalent:

$$H_t = \mathrm{E}_t\left[W_T - \int_t^T \left\{\frac{1}{2}\gamma_1||\sigma_{H,s}||^2 + \frac{1}{2}\gamma_2||\sigma_{H,x}||^2\right\} du\right]$$

which is equation (7) in the main text with $\gamma_1 = \gamma$ and $\gamma_2 = \gamma_x$.

For $\gamma_x \neq \gamma$, this is the certainty equivalent of a source-dependent stochastic differential utility agent as advocated in Skiadas (2008). We also give a recursive heuristic argument for the construction of this certainty equivalent (following Skiadas (2008)) in the following section.

If we pick $\gamma_1 = \gamma_2 = \gamma$, then our derivation in Appendix A implies that H_t is the certainty equivalent of a CARA agent with absolute risk-aversion coefficient γ. That is following the derivation in the previous section we obtain:

$$\text{If } \gamma_1 = \gamma_2 = \gamma \text{ then } H_t = -\frac{1}{\gamma}\log \mathrm{E}_t[e^{-\gamma W_T}]$$

In general, with W_t dynamics given in (6) above, we look for a solution of the form $H_t = W_t + J(x_t, n_t, t)$. Plugging this guess into the BSDE, we find $J(x_t, n_t, t)$ satisfies (note that this guess

[17] Following Skiadas (2008) one can also give a heuristic derivation of this recursive utility based on a specific source-dependent certainty equivalent as in Appendix F.

also implies that the diffusion of H has two components $\sigma_{H,s} = n_t^\top \sigma_s + J_x^\top \sigma_{xs}$ and $\sigma_{H,x} = J_x^\top \sigma_x$):

$$J(x_t, n_t, t) = \mathrm{E}_t \left[\int_t^T \{ dW_u - \frac{1}{2}\gamma_1 n_u^\top \sigma_s \sigma_s^\top n_u du - \frac{1}{2} J_x^\top (\gamma_1 \sigma_{xs} \sigma_{xs}^\top + \gamma_2 \sigma_x \sigma_x^\top) J_x du - \gamma_1 n_u^\top \sigma_s \sigma_{xs}^\top J_x du \right]$$

which is, indeed, the objective function we consider in equations (102) and (110) below with $\gamma_1 = \gamma$ and $\gamma_2 = \gamma_x$.

Now, we also see that if $\gamma_2 \sigma_x = 0$ and $\sigma_{xs} = 0$, then the certainty equivalent indeed reduces to the CMV objective function as claimed in Theorem 2, that is (with $\gamma_1 = \gamma$):

$$J(x_t, t) = \mathrm{E}_t \left[\int_t^T dW_u - \frac{1}{2}\gamma dW_u^2 \right]$$

C Recursive Construction of the 'Source-Dependent' Stochastic Differential Utility of Terminal Wealth

Following Skiadas (2008) and Hugonnier, Pelgrin, and St-Amour (2012), we consider a local approximation argument to show heuristically how to construct recursively the certainty equivalent H_t of our agent who consumes only at maturity T and has source-dependent risk-aversion. We assume wealth is driven by two independent Brownian motions Z^x, Z^s and one Poisson jump N_t with an arrival intensity of ρ. We allow for a jump to deal with the possible random horizon model. We also assume that prior to t, the certainty equivalent has dynamics given by:

$$dH_t = \mu_H dt + \sigma_{H,s} dZ^s + \sigma_{H,x} dZ^x + \eta_H (dN_t - \rho dt). \tag{96}$$

At any time $t < T$ the certainty equivalent is defined by the following recursion

$$\mathcal{U}(H_t, 0, 0, 0) = \mathrm{E}_t[\mathcal{U}(H_t + \mu_H dt, \sigma_{H,s} dZ^s, \sigma_{H,x} dZ^x, \eta_H (dN_t - \rho dt))] \tag{97}$$

with the boundary condition $H_T = W_T$, for some source-dependent risk-aversion function $\mathcal{U}(z_0, z_1, z_2, z_3)$. Note that if $\mathcal{U}(z_0, z_1, z_2, z_3) = U(z_0 + z_1 + z_2 + z_3)$ we obtain the same recursive definition as in the section B. Instead, here we assume the following function:

$$\mathcal{U}(z_0, z_1, z_2, z_3) = U_1(z_0 + z_1) + \frac{U_1'(z_0)}{U_2'(z_0)}(U_2(z_0 + z_2) - U_2(z_0)) + \frac{U_1'(z_0)}{U_3'(z_0)}(U_3(z_0 + z_3) - U_3(z_0))$$

Using this we can rewrite the recursion (97), using the Itô rule for the right-hand side as:

$$\begin{aligned}U_1(H_t) =& U_1(H_t) + U_1'(H_t)\mu_H dt + \frac{1}{2}U_1''(H_t)\sigma_{H,s}^2 dt + \frac{U_1'(H_t)}{U_2'(H_t)}\frac{1}{2}U_2''(H_t)\sigma_{H,x}^2 dt\\ &- U_1'(H_t)\left(\eta_H - \frac{U_3(H_t+\eta_H) - U_3(H_t)}{U_3'(H_t)}\right)\rho dt\end{aligned}$$

Simplifying and rewriting we obtain the driver μ_H of the BSDE which defines the source-dependent SDU:

$$\mu_H = -\frac{1}{2}\frac{U_1''(H)}{U_1'(H)}||\sigma_{H,s}||^2 - \frac{1}{2}\frac{U_2''(H)}{U_2'(H)}||\sigma_{H,x}||^2 + \rho\left(\eta_H - \frac{U_3(H_t+\eta_H) - U_3(H_t)}{U_3'(H_t)}\right) \tag{98}$$

If we specialize to CARA utility functions $U_i(x) = -e^{-\gamma_i x}$, then the BSDE representation becomes

$$H_t = \mathrm{E}_t\left[W_T - \int_t^T \left\{\frac{1}{2}\gamma_1||\sigma_{H,s}||^2 + \frac{1}{2}\gamma_2||\sigma_{H,x}||^2 + \rho(\eta_H - \frac{1-e^{-\gamma_3\eta_H}}{\gamma_3})\right\}\right] \tag{99}$$

When there are no jumps (i.e., $\rho = 0$) then this is the driver of the BSDE corresponding to recursive preferences with source-dependent risk aversion that we introduced in (94). The jump component is useful to understand the stationary case where the horizon is generated by the first jump of the poisson process.

D Source-dependent SDU with a random horizon

We consider the generalization of our SDU definition where $\mathcal{T}$ is generated by the first jump of a Poisson process with intensity ρ.

Then we define the certainty equivalent as the solution $(H_t, \sigma_{H,s}, \sigma_{H,x}, \eta_H := W_t - H_{t^-})$ to the recursive BSDE defined for $t \leq \mathcal{T}$:

$$\begin{aligned}H_t =& \mathrm{E}_t\left[W_{\mathcal{T}} - \int_t^{\mathcal{T}} \left\{\frac{1}{2}\gamma_1||\sigma_{H,s}||^2 + \frac{1}{2}\gamma_2||\sigma_{H,x}||^2 + \rho\left(W_s - H_{s^-} - \frac{1-e^{-\gamma_3(W_s - H_{s^-})}}{\gamma_3}\right)\right\} ds\right]\\ =& W_t\\ &+ \mathbf{1}_{\{\mathcal{T}>t\}}\mathrm{E}_t\left[\int_t^\infty e^{-\rho(s-t)}\left\{dW_s - \left[\frac{1}{2}\gamma_1||\sigma_{H,s}||^2 + \frac{1}{2}\gamma_2||\sigma_{H,x}||^2 + \rho\left(W_s - H_{s^-} - \frac{1-e^{-\gamma_3(W_s-H_{s^-})}}{\gamma_3}\right)\right] ds\right\}\right]\end{aligned} \tag{100}$$

The equality between the first and second line requires an additional transversality condition.[18]

We prove our first result.

Theorem 10 *When $\gamma_1 = \gamma_2 = \gamma_3 = \gamma$ then the solution to (100) is the certainty equivalent of a*

[18] Note that $\mathrm{E}_t\left[\int_t^{\mathcal{T}} dX_u\right] = \mathrm{E}_t\left[\int_t^\infty \rho e^{-\rho(s-t)}ds \int_t^s dX_u\right] = \mathrm{E}_t\left[\int_t^\infty e^{-\rho(s-t)}dX_s - [e^{-\rho(s-t)}(X_s - X_t)]_t^\infty\right]$. Therefore the transversality condition is $\lim_{T\to\infty}\mathrm{E}[e^{-\rho T}X_T] = 0$.

CARA investor with expected utility of terminal wealth generated at the random horizon $\mathcal{T}$. That is $H_t = \frac{1}{\gamma}\log(\mathrm{E}_t[e^{-\gamma W_{\mathcal{T}}}])$.

Proof. Note that the solution to (100) when $\gamma_1 = \gamma_2 = \gamma_3 = \gamma$ is a jump diffusion process, with the property that $H_{\mathcal{T}} = W_{\mathcal{T}}$ at the jump time. Therefore we posit the following dynamics for H_t on $\mathcal{T} > t$:

$$dH_t = \mu_H dt + \sigma_{H,s} dZ^s + \sigma_{H,x} dZ^x + (W_t - H_{t^-})(d\mathbf{1}_{\{\mathcal{T}\le t\}} - \rho dt) \tag{101}$$

From the BSDE definition we can see that the drift μ_H (on $\tau > t$) is given by:

$$\mu_H = \left\{ \frac{1}{2}\gamma||\sigma_{H,s}||^2 + \frac{1}{2}\gamma||\sigma_{H,x}||^2 + \rho\left(W_s - H_{s^-} - \frac{1 - e^{-\gamma(W_s - H_{s^-})}}{\gamma}\right)\right\}$$

Applying Itô's lemma we find $U(H_t) = -e^{-\gamma H_t}$ has dynamics:

$$\begin{aligned} dU(H_t) =& \{-\frac{1}{2}U''(H)(||\sigma_{H,x}||^2 + ||\sigma_{H,s}||^2) + U'(H_{t^-})(\mu_H - \rho(W_t - H_{t^-})\}dt \\ &+ U'(H)\sigma_{H,s}dZ^s + U'(H)\sigma_{H,x}dZ^x + (U(W_t) - U(H_{t^-}))d\mathbf{1}_{\{\mathcal{T}\le t\}} \\ =& U'(H)\sigma_{H,s}dZ^s + U'(H)\sigma_{H,x}dZ^x + (U(W_t) - U(H_{t^-}))(d\mathbf{1}_{\{\mathcal{T}\le t\}} - \rho dt) \end{aligned}$$

where we have substituted the expression for μ_H to get the second equality.

Therefore we find that the solution to the BSDE is such that $U(H_t)$ is martingale, which takes on the value $u(W_{\mathcal{T}})$ at $\mathcal{T}$. I follows that at $t < \mathcal{T}$ and using the optional stopping theorem:

$$U(H_t) = \mathrm{E}_t[U(H_{\mathcal{T}})] = \mathrm{E}_t[U(W_{\mathcal{T}})]$$

which is the desired result. ■

Note that this investor has same risk-aversion to the three types of shocks $Z^s, Z^x, \mathcal{T}$.

E Finite horizon solution without transaction costs

Without transaction costs (i.e., when $\Lambda = 0$), we optimize directly over the number of shares n_t as the wealth-dynamics simplifies and the optimal trading will have infinite variation. We look for a solution of the form $H_t = W_t + J(x_t, t)$, which implies $\sigma_{H,s} = n^\top\sigma_s + J_x^\top\sigma_{xs}$ and $\sigma_{H,x} = J_x^\top\sigma_x$. It follows from equation (7) that the function $J(x,t)$ must satisfy:

$$J(x_t, t) = \max_n \mathrm{E}_t\left[\int_t^T \left\{dW_u - \frac{1}{2}\gamma n_u^\top\Sigma n_u du - \frac{1}{2}J_x^\top\Omega J_x du - \gamma n_u^\top\Sigma_{sx}J_x du\right\}\right] \tag{102}$$

where we define:

$$\Omega = \gamma\sigma_{xs}\sigma_{xs}^\top + \gamma_x\sigma_x\sigma_x^\top \tag{103}$$

$$\Sigma_{sx} = \sigma_s\sigma_{xs}^\top \tag{104}$$

The corresponding Bellman-equation is:

$$0 = \max_n \mathrm{E}_t[dW_t - \frac{1}{2}\gamma n_t^\top \Sigma n_t dt - \frac{1}{2}J_x^\top \Omega J_x dt - \gamma n_t^\top \Sigma_{sx} J_x dt + dJ(t, x_t)] \tag{105}$$

Using the definition of the wealth equation (with $\Lambda = 0$) we obtain

$$0 = \max_n \left\{ n^\top(\mu_0 + \mu x) - \frac{1}{2}\gamma n^\top \Sigma n - \frac{1}{2}J_x^\top \Omega J_x - \gamma n^\top \Sigma_{sx} J_x + J_t - J_x^\top \kappa x + \frac{1}{2}\,\mathrm{Tr}(J_{xx}\Sigma_x) \right\}$$

and we have defined J_x and J_{xx} as respectively the gradient and hessian of $J(x,t)$ with respect to x, and J_t the partial derivative with respect to t.

The first order condition, with respect to n, is

$$n = (\gamma\Sigma)^{-1}(\mu_0 + \mu x - \gamma\Sigma_{sx}J_x)$$

Plugging back into the HJB equation we get:

$$0 = \frac{1}{2}(\mu_0 + \mu x - \gamma\Sigma_{sx}J_x)^\top (\gamma\Sigma)^{-1}(\mu_0 + \mu x - \gamma\Sigma_{sx}J_x) - \frac{1}{2}J_x^\top \Omega J_x + J_t - J_x^\top \kappa x + \frac{1}{2}\,\mathrm{Tr}(J_{xx}\Sigma_x)$$

We guess that the value function is of the form:

$$J(x,t) = c_0(t) + c_1(t)^\top x + \frac{1}{2}x^\top c_2(t) x$$

where c_2 is symmetric (w.l.og.) and c_0, c_1 are all matrices (with appropriate dimensions) of deterministic functions.

$$\begin{aligned} J_t &= \dot{c_0} + \dot{c_1}^\top x + \frac{1}{2}x^\top \dot{c_2} x \\ J_x &= c_1 + c_2 x \\ J_{xx} &= c_2 \end{aligned}$$

Thus HJB becomes

$$-\dot{c_0} - \dot{c_1}^\top x - \frac{1}{2}x^\top \dot{c_2}x = \frac{1}{2}\left(\mu_0 + \mu x - \gamma\Sigma_{sx}(c_1 + c_2 x)\right)^\top (\gamma\Sigma)^{-1}\left(\mu_0 + \mu x - \gamma\Sigma_{sx}(c_1 + c_2 x)\right)$$
$$- \frac{1}{2}(c_1 + c_2 x)^\top \Omega(c_1 + c_2 x) - (c_1 + c_2 x)^\top \kappa x + \frac{1}{2}\operatorname{Tr}(c_2\Sigma_x)$$

This equation is satisfied provided c_0, c_1, c_2 solve the following system:

$$\begin{aligned}
-\dot{c_0} &= \frac{1}{2}\left(\mu_0 - \gamma\Sigma_{sx}c_1\right)^\top (\gamma\Sigma)^{-1}\left(\mu_0 - \gamma\Sigma_{sx}c_1\right) - \frac{1}{2}c_1^\top \Omega c_1 + \frac{1}{2}\operatorname{Tr}(c_2\Sigma_x) \\
-\dot{c_1} &= \left(\mu - \gamma\Sigma_{sx}c_2\right)^\top (\gamma\Sigma)^{-1}\left(\mu_0 - \gamma\Sigma_{sx}c_1\right) - c_2^\top \Omega c_1 - \kappa^\top c_1 \\
-\dot{c_2} &= \left(\mu - \gamma\Sigma_{sx}c_2\right)^\top (\gamma\Sigma)^{-1}\left(\mu - \gamma\Sigma_{sx}c_2\right) - c_2^\top \Omega c_2 - 2c_2^\top \kappa
\end{aligned} \tag{106}$$

This system has to be solved subject to the boundary condition $c_0(T) = 0$, $c_1(T) = 0$ and $c_2(T) = 0$ (where 0 is the matrix of zeros with appropriate dimension).

We note that the if $\mu_0 = 0$ then $c_1(t) = 0$ and the trading strategy only depends on c_2 which solves an autonomous ODE of the Riccatti type:

$$-\dot{c_1} = \left(\mu - \gamma\Sigma_{sx}c_2\right)^\top (\gamma\Sigma)^{-1}\mu_0 - \{(\mu - \gamma\Sigma_{sx}c_2)^\top \Sigma^{-1}\Sigma_{sx} + c_2^\top \Omega + \kappa^\top\}c_1 \tag{107}$$

$$-\dot{c_2} = c_2^\top \left(\gamma\Sigma_{sx}^\top \Sigma^{-1}\Sigma_{sx} - \Omega\right) c_2 + 2c_2^\top(-\kappa - \Sigma_{sx}\Sigma^{-1}\mu) + \mu^\top(\gamma\Sigma)^{-1}\mu \tag{108}$$

The solution is easily obtained numerically. In terms of the solution the optimal position is given by:

$$n_t = (\gamma\Sigma)^{-1}(\mu_0 + \mu x_t) - \Sigma^{-1}\Sigma_{sx}(c_1(t) + c_2(t)x)$$

where we see that it can be decomposed into the CMVE Markowitz portfolio and a hedging portfolio (Merton (1973)). In the absence of transaction costs the investor will choose to deviate from the Markowitz portfolio as soon as $\Sigma_{sx} \neq 0$.

In particular, we note that the GP investor (who effectively acts as if $\Sigma_{sx} = 0$ and with $\gamma_x = 0$, see Remark 3) is myopic in the sense that, absent transaction costs (i.e., if $\Lambda = 0$), she would choose to hold the CMVE instantaneous mean-variance efficient Markowitz portfolio at all times:

$$CMVE_t = (\gamma\Sigma)^{-1}(\mu_0 + \mu x_t) \tag{109}$$

Of course, with transaction costs the optimal portfolio will deviate from the Markowitz portfolio both for the GP investor and the non-myopic CARA agent. We now turn to the case with transaction costs.

F Finite horizon solution with transaction costs

We now consider the case with transaction costs when $\Lambda \neq 0$. We look for a solution of the form $H_t = W_t + J(n_t, x_t, t)$, which implies $\sigma_{H,s} = n^\top \sigma_s + J_x^\top \sigma_{xs}$ and $\sigma_{H,x} = J_x^\top \sigma_x$. It follows that the function $J(n, x, t)$ must satifsy:

$$J(n_t, x_t, t) = \max_\theta \mathrm{E}_t \left[\int_t^T \left\{ dW_u - \frac{1}{2}\gamma n_u^\top \Sigma n_u du - \frac{1}{2} J_x^\top \Omega J_x du - \gamma n_u^\top \Sigma_{sx} J_x du \right\} \right] \tag{110}$$

where we define:

$$\Omega = \gamma \sigma_{xs} \sigma_{xs}^\top + \gamma_x \sigma_x \sigma_x^\top \tag{111}$$

$$\Sigma_{sx} = \sigma_s \sigma_{xs}^\top \tag{112}$$

Thus $J(n, x, t)$ satisfies the HJB equation:

$$0 = \max_\theta \mathrm{E}_t [dW_t - \frac{1}{2}\gamma n_t^\top \Sigma n_t dt - \frac{1}{2} J_x^\top \Omega J_x dt - \gamma n_t^\top \Sigma_{sx} J_x dt + dJ(t, n_t, x_t)] \tag{113}$$

Using the dynamics of the wealth process, we obtain the following equation:

$$0 = \max_\theta \left\{ n^\top (\mu_0 + \mu x) - \frac{1}{2}\theta^\top \Lambda \theta - \frac{1}{2}\gamma n^\top \Sigma n - \frac{1}{2} J_x^\top \Omega J_x - \gamma n^\top \Sigma_{sx} J_x + J_t + J_n^\top \theta - J_x^\top \kappa x + \frac{1}{2} \mathrm{Tr}(J_{xx} \Sigma_x) \right\}$$

and we have defined J_x and J_{xx} as respectively the gradient and hessian of $J(n, x, t)$ with respect to x, J_n the gradiant with respect to n, and J_t the partial derivative with respect to t.

The first order condition is

$$\theta = \Lambda^{-1} J_n$$

Plugging back into the HJB equation we get:

$$0 = \max_\theta \left\{ n^\top (\mu_0 + \mu x) + \frac{1}{2} J_n^\top \Lambda^{-1} J_n - \frac{1}{2}\gamma n^\top \Sigma n - \frac{1}{2} J_x^\top \Omega J_x - \gamma n^\top \Sigma_{sx} J_x + J_t - J_x^\top \kappa x + \frac{1}{2} \mathrm{Tr}(J_{xx} \Sigma_x) \right\}$$

We guess that the value function is of the form:

$$J(n, x, t) = -\frac{1}{2} n^\top Q(t) n + n^\top (q_0(t) + q(t)^\top x) + c_0(t) + c_1(t)^\top x + \frac{1}{2} x^\top c_2(t) x$$

where Q, c_2 are symmetric (w.l.og.) square matrices and q_0, q_1, c_0, c_1 are all matrices or vectors (with appropriate dimensions) of deterministic functions.

$$
\begin{aligned}
J_t &= -\frac{1}{2}n^\top \dot{Q} n + n^\top(\dot{q_0} + \dot{q}^\top x) + \dot{c_0} + \dot{c_1}^\top x + \frac{1}{2}x^\top \dot{c_2} x \\
J_n &= -Qn + q_0 + q^\top x \\
J_x &= qn + c_1 + c_2 x \\
J_{xx} &= c_2
\end{aligned}
$$

Thus HJB becomes

$$
\begin{aligned}
0 = &-\frac{1}{2}n^\top \dot{Q} n + n^\top(\dot{q_0} + \dot{q}^\top x) + \dot{c_0} + \dot{c_1}^\top x + \frac{1}{2}x^\top \dot{c_2} x \\
&+ \frac{1}{2}(-Qn + q_0 + q^\top x)^\top \Lambda^{-1}(-Qn + q_0 + q^\top x) + n^\top(\mu_0 + \mu x) - \frac{1}{2}\gamma n^\top \Sigma n \\
&- \frac{1}{2}(qn + c_1 + c_2^\top x)^\top \Omega (qn + c_1 + c_2 x) - \gamma n^\top \Sigma_{sx}(qn + c_1 + c_2 x) - (qn + c_1 + c_2 x)^\top \kappa x + \frac{1}{2}\operatorname{Tr}(c_2 \Sigma_x)
\end{aligned}
$$

Rewriting:

$$
\begin{aligned}
0 = &\frac{1}{2}n^\top(-\dot{Q} + Q\Lambda^{-1}Q - \gamma\Sigma - q^\top \Omega q - 2\gamma \Sigma_{sx} q)n \\
&+ n^\top(\dot{q_0} + \mu_0 - Q\Lambda^{-1}q_0 - q^\top \Omega c_1 - \gamma \Sigma_{sx} c_1) + n^\top(\dot{q}^\top - Q\Lambda^{-1}q^\top + \mu - q^\top \kappa - q^\top \Omega c_2 - \gamma \Sigma_{sx} c_2)x \\
&+ x^\top(\frac{1}{2}\dot{c_2} + \frac{1}{2}q\Lambda^{-1}q^\top - c_2\kappa - \frac{1}{2}c_2 \Omega c_2)x \\
&+ x^\top(\dot{c_1} + q\Lambda^{-1}q_0 - c_2 \Omega c_1 - \kappa^\top c_1) + \dot{c_0} + \frac{1}{2}q_0^\top \Lambda^{-1} q_0 - \frac{1}{2}c_1^\top \Omega c_1 + \frac{1}{2}\operatorname{Tr}(c_2 \Sigma_x)
\end{aligned}
$$

So we obtain the set of ODEs that need to be satisfied by the solution.

$$
-\dot{Q} = \gamma\Sigma - Q\Lambda^{-1}Q + q^\top \Omega q + \gamma(\Sigma_{sx} q + q^\top \Sigma_{sx}^\top) \tag{114}
$$

$$
-\dot{q}^\top = \mu - q^\top \kappa - Q\Lambda^{-1}q^\top - q^\top \Omega c_2 - \gamma \Sigma_{sx} c_2 \tag{115}
$$

$$
-\dot{c_2} = -(c_2\kappa + \kappa^\top c_2) + q\Lambda^{-1}q^\top - c_2 \Omega c_2 \tag{116}
$$

$$
-\dot{q_0} = \mu_0 - Q\Lambda^{-1}q_0 - q^\top \Omega c_1 - \gamma \Sigma_{sx} c_1 \tag{117}
$$

$$
-\dot{c_1} = -\kappa^\top c_1 + q\Lambda^{-1}q_0 - c_2 \Omega c_1 \tag{118}
$$

$$
-\dot{c_0} = \frac{1}{2}\operatorname{Tr}(c_2 \Sigma_x) + \frac{1}{2}q_0^\top \Lambda^{-1} q_0 - \frac{1}{2}c_1^\top \Omega c_1 \tag{119}
$$

subject to boundary conditions $Q(T) = 0$, $q(T) = 0$, $q_0(T) = 0$, $c_0(T) = 0$, $c_1(T) = 0$, and $c_2(T) = 0$. We note that if $\mu_0 = 0$ then $c_1(t) = 0$ and $q_0(t) = 0$, $\forall t$.

Also, if $\Omega = 0$ (for example in the GP case, where there is no correlation $\Sigma_{xs} = 0$ and there is vanishing risk-aversion to Z^s risk, that is $\gamma_x = 0$) then the system for Q, q is autonomous and does not depend on the solution for c_2, whereas when there is a hedging demand $\gamma_x > 0$ then the system for Q, q, c_2 needs to be solved jointly. So c_2 encodes the hedging demand component, just like in the case without transaction costs.

To interpret the optimal trading strategy, note that the value function is maximized with respect to the position vector n at the optimal aim portfolio:

$$aim(x_t, t) = Q^{-1}(t)(q_0(t) + q(t)^\top x_t).$$

Since $J_n = -Qn + q_0 + q^\top x$ the optimal trade can be written as:

$$\theta = \Lambda^{-1} J_n = \Lambda^{-1} Q(aim(x_t, t) - n_t)$$

So with the definition of trade intensity $\tau_t = \Lambda^{-1} Q(t)$ we get the optimal trading strategy:

$$dn_t = \tau_t(aim(x_t, t) - n_t)dt \tag{120}$$

G The finite horizon solution CMV preferences (i.e., $\sigma_{xs} = 0$ and $\gamma_x = 0$)

As discussed in Remark 2, the solution to the finite horizon model where agents have CMV preferences (as in equation (10)) corresponds to the solution of the source-dependent risk-aversion recursive utility agent with parameters restricted to $\sigma_{xs} = 0$ and $\gamma_x = 0$ (which implies $\Omega = 0$). To understand the optimal trading rule (aim_t, τ_t) the relevant system of ODE we need to solve becomes:

$$-\dot{Q} = \gamma\Sigma - Q\Lambda^{-1}Q \tag{121}$$

$$-\dot{q}^\top = \mu - q^\top \kappa - Q\Lambda^{-1} q^\top \tag{122}$$

$$-\dot{q_0} = \mu_0 - Q\Lambda^{-1} q_0 \tag{123}$$

Now, we can rewrite this system in terms of the trading speed matrix $\tau = \Lambda^{-1}Q$ as:

$$-\dot{\tau} = \gamma\Lambda^{-1}\Sigma - \tau\tau \tag{124}$$

$$-\dot{Q} = \gamma\Sigma - \tau^\top Q \tag{125}$$

$$-\dot{q}^\top = \mu - q^\top \kappa - \tau^\top q^\top \tag{126}$$

$$-\dot{q_0} = \mu_0 - \tau^\top q_0 \tag{127}$$

This system has an intuitive closed-form solution. Let us define the diagonalization $\gamma\Lambda^{-1}\Sigma = FD_\eta F^\top$, where we define D_η as the diagonal matrix with eigenvalue η_i on the i^{th} diagonal. Then we see that $\tau = FD_h F^\top$ where D_h is the diagonal matrix with the deterministic function $h_i(t)$ on its i^{th} diagonal. Plugging into the ODE for τ we find that the solution separates into n individual

ODEs for the h_i functions, which solve:

$$-\dot{h}_i = \eta_i - h_i^2 \quad s.t. \;\; h_i(T) = 0 \tag{128}$$

The solution is as given in the theorem.

It follows that the trading speed matrix is given by $\tau_t = F D_h(t) F^\top$ and the Q matrix is $Q(t) = \Lambda \tau(t)$.

To solve for $q(t), q_0(t)$, we use the following lemma:

lemma 1 *The following holds:*

$$\frac{d}{dt} e^{-\int_0^t \tau_s^\top ds} = -\tau_t^\top e^{-\int_0^t \tau_s^\top ds} dt = -e^{-\int_0^t \tau_s^\top ds} \tau_t^\top dt \tag{129}$$

Further, $\forall t, u, T$ *the following holds:*

$$e^{-\int_t^u \tau_s^\top ds} e^{-\int_u^T \tau_s^\top ds} = e^{-\int_t^T \tau_s^\top ds} \tag{130}$$

Proof. Note that $\int_0^t \tau_s^\top ds = F D_{\int_0^t h(s) ds} F^\top$. Therefore, from the properties of the matrix exponential[19] $e^{-\int_0^t \tau_s^\top ds} = F D_{e^{-\int_0^t h(s)ds}} F^\top$. Now, taking the derivative we find:

$$\frac{d}{dt} e^{-\int_0^t \tau_s^\top ds} = F D_{e^{-\int_0^t h(s)ds} h(t)} F^\top \tag{131}$$

$$= F D_{e^{-\int_0^t h(s)ds}} F^\top F D_{-h(t)} F^\top \tag{132}$$

$$= e^{-\int_0^t \tau_s^\top ds} \tau_t \tag{133}$$

which proves the first equality of the first statement. The second equality of the first statement follows immediately from using the fact that two diagonal matrices commute in the second line above.

Now to prove the second statement, we proceed similarly to above and note:

$$e^{-\int_t^u \tau_s^\top ds} e^{-\int_u^T \tau_s^\top ds} = F D_{e^{-\int_t^u h_s ds}} F^\top F D_{e^{-\int_u^T h_s ds}} F^\top \tag{134}$$

$$= F D_{e^{-\int_t^T h_s ds}} F^\top \tag{135}$$

$$= e^{-\int_t^T \tau_s^\top ds} \tag{136}$$

■

Now, we can use this lemma to solve the ODE system. We find:

[19] The matrix exponential is defined as $e^X = \sum_{k=0}^{\infty} \frac{X^k}{k!}$. It follows that if $XY = YX$ then $e^{X+Y} = e^X e^Y$. Further, if Y is invertible then $e^{YXY^{-1}} = Y e^X Y^{-1}$. Recall that $FF^\top = I$.

lemma 2 *The solution to the ODE system is given as follows;*

$$q(t) = \int_t^T e^{-\int_t^u \tau_s^\top ds} \mu e^{-\int_t^u \kappa ds} du \tag{137}$$

$$q_0(t) = \int_t^T e^{-\int_t^u \tau_s^\top ds} du \mu_0 \tag{138}$$

$$Q(t) = \int_t^T e^{-\int_t^u \tau_s^\top ds} du \gamma \Sigma \tag{139}$$

Proof. We prove only the first results as the other ones are proved similarly. Using lemma 1 we have:

$$\frac{d}{dt} e^{\int_0^t \tau_s^\top ds} q(t)^\top e^{\int_0^t \kappa ds} = e^{\int_0^t \tau_s^\top ds} (\tau_t^\top q(t)^\top + \dot{q(t)}^\top + q(t)^\top \kappa) e^{\int_0^t \kappa ds} \tag{140}$$

$$= e^{\int_0^t \tau_s^\top ds} \mu e^{\int_0^t \kappa ds} \tag{141}$$

Now integrating and using the boundary condition $q(T) = 0$ we find:

$$e^{\int_0^t \tau_s^\top ds} q(t)^\top e^{\int_0^t \kappa ds} = -\int_t^T e^{\int_0^u \tau_s^\top ds} \mu e^{\int_0^u \kappa ds} du \tag{142}$$

Now left-multiplying by $e^{\int_t^0 \tau_s^\top ds}$ and right-multiplying by $e^{\int_t^0 \kappa ds}$ and using lemma 1 we find the desired expression. ■

The main result then follows from the definition of the aim portfolio $aim(t, x) = Q(t)^{-1}(q_0(t) + q(t)x)$.

H The infinite horizon portfolio problem without transaction costs ($\Lambda = 0$)

Without transaction costs (i.e., when $\Lambda = 0$), we optimize directly over the number of shares n_t as the wealth-dynamics simplifies and the optimal trading will have infinite variation. Different from the finite horizon case, we now look for a stationary solution of the form $H_t = W_t + J(x_t)$, which implies $\sigma_{H,s} = n^\top \sigma_s + J_x^\top \sigma_{xs}$ and $\sigma_{H,x} = J_x^\top \sigma_x$. It follows from equation (32) that the function $J(x)$ must satisfy:

$$J(x_t) = \max_n \mathrm{E}_t \left[\int_t^\infty e^{-\rho(u-t)} \left\{ dW_u - \frac{1}{2}\gamma n_u^\top \Sigma n_u du - \frac{1}{2} J_x^\top \Omega J_x du - \gamma n_u^\top \Sigma_{sx} J_x du \right\} \right] \tag{143}$$

The corresponding Bellman-equation is:

$$0 = \max_n \mathrm{E}_t [dW_t - \frac{1}{2}\gamma n_t^\top \Sigma n_t dt - \frac{1}{2} J_x^\top \Omega J_x dt - \gamma n_t^\top \Sigma_{sx} J_x dt + dJ(x_t) - \rho J(x_t)] \tag{144}$$

Using the definition of the wealth equation (with $\Lambda = 0$) we obtain

$$\rho J(x_t) = \max_n \left\{ n^\top(\mu_0 + \mu x) - \frac{1}{2}\gamma n^\top \Sigma n - \frac{1}{2}J_x^\top \Omega J_x - \gamma n^\top \Sigma_{sx} J_x - J_x^\top \kappa x + \frac{1}{2}\operatorname{Tr}(J_{xx}\Sigma_x) \right\}$$

and we have defined J_x and J_{xx} as respectively the gradient and hessian of $J(x)$ with respect to x.

The first order condition, with respect to n, is

$$n = (\gamma\Sigma)^{-1}(\mu_0 + \mu x - \gamma \Sigma_{sx} J_x)$$

Plugging back into the HJB equation we get:

$$\rho J = \frac{1}{2}(\mu_0 + \mu x - \gamma\Sigma_{sx}J_x)^\top (\gamma\Sigma)^{-1}(\mu_0 + \mu x - \gamma\Sigma_{sx}J_x) - \frac{1}{2}J_x^\top \Omega J_x - J_x^\top \kappa x + \frac{1}{2}\operatorname{Tr}(J_{xx}\Sigma_x)$$

We guess that the value function is of the form:

$$J(x) = c_0 + c_1^\top x + \frac{1}{2}x^\top c_2 x$$

where c_2 is symmetric (w.l.og.) and c_0, c_1 are all matrices (with appropriate dimensions) of constants.

$$\begin{aligned} J_x &= c_1 + c_2 x \\ J_{xx} &= c_2 \end{aligned}$$

Thus the HJB equation becomes

$$\begin{aligned} \rho(c_0 + c_1^\top x + \frac{1}{2}x^\top c_2 x) =& \frac{1}{2}(\mu_0 + \mu x - \gamma\Sigma_{sx}(c_1 + c_2 x))^\top (\gamma\Sigma)^{-1}(\mu_0 + \mu x - \gamma\Sigma_{sx}(c_1 + c_2 x)) \\ &- \frac{1}{2}(c_1 + c_2 x)^\top \Omega (c_1 + c_2 x) - (c_1 + c_2 x)^\top \kappa x + \frac{1}{2}\operatorname{Tr}(c_2\Sigma_x) \end{aligned}$$

This equation is satisfied provided c_0, c_1, c_2 solve the following system:

$$\begin{aligned} \rho c_0 &= \frac{1}{2}(\mu_0 - \gamma\Sigma_{sx}c_1)^\top (\gamma\Sigma)^{-1}(\mu_0 - \gamma\Sigma_{sx}c_1) - \frac{1}{2}c_1^\top \Omega c_1 + \frac{1}{2}\operatorname{Tr}(c_2\Sigma_x) \qquad (145) \\ \rho c_1 &= (\mu - \gamma\Sigma_{sx}c_2)^\top (\gamma\Sigma)^{-1}(\mu_0 - \gamma\Sigma_{sx}c_1) - c_2^\top \Omega c_1 - \kappa^\top c_1 \\ \rho c_2 &= (\mu - \gamma\Sigma_{sx}c_2)^\top (\gamma\Sigma)^{-1}(\mu - \gamma\Sigma_{sx}c_2) - c_2^\top \Omega c_2 - 2c_2^\top \kappa \end{aligned}$$

In addition we require that the transversality condition $\lim_{T\to\infty} \mathrm{E}[e^{-\rho T}J(X_T)] = 0$ be satisfied.[20]

We note that the if $\mu_0 = 0$ then $c_1 = 0$ and the trading strategy only depends on c_2 which solves an autonomous ODE of the Riccatti type:

$$0 = (\mu - \gamma\Sigma_{sx}c_2)^\top (\gamma\Sigma)^{-1}\mu_0 - \{(\mu - \gamma\Sigma_{sx}c_2)^\top\Sigma^{-1}\Sigma_{sx} + c_2^\top\Omega + \kappa^\top + \rho\}c_1 \tag{146}$$

$$0 = c_2\left(\gamma\Sigma_{sx}^\top\Sigma^{-1}\Sigma_{sx} - \Omega\right)c_2 - c_2(\rho + 2\kappa + 2\Sigma_{sx}\Sigma^{-1}\mu) + \mu^\top(\gamma\Sigma)^{-1}\mu \tag{147}$$

The solution is easily obtained numerically. In terms of the solution the optimal position is given by:

$$n_t = (\gamma\Sigma)^{-1}(\mu_0 + \mu x_t) - \Sigma^{-1}\Sigma_{sx}(c_1 + c_2 x)$$

where we see that it can be decomposed into the CMVE Markowitz portfolio and a hedging portfolio (Merton (1973)). In the absence of transaction costs the investor will choose to deviate from the Markowitz portfolio as soon as $\Sigma_{sx} \neq 0$.

In particular, we note that, as in the finite-horizon case, the GP investor (who effectively acts as if $\Sigma_{sx} = 0$ and with $\gamma_x = 0$, see Remark 3) is myopic in the sense that, absent transaction costs (i.e., if $\Lambda = 0$), she would choose to hold the CMVE Markowitz portfolio at all times:

$$Mwz_t = (\gamma\Sigma)^{-1}(\mu_0 + \mu x_t) \tag{148}$$

Of course, with transaction costs the optimal portfolio will deviate from the Markowitz portfolio both for the GP investor and the non-myopic CARA agent. We now turn to the infinite horizon case with transaction costs.

I The infinite horizon portfolio problem with transaction costs ($\Lambda \neq 0$)

We now consider the optimal portfolio choice of a source-dependent utility agent with objective function (32) for the case with transaction costs when $\Lambda \neq 0$. We look for a solution of the form $H_t = W_t + J(n_t, x_t)$, which implies $\sigma_{H,s} = n^\top\sigma_s + J_x^\top\sigma_{xs}$ and $\sigma_{H,x} = J_x^\top\sigma_x$. It follows that the

[20] Indeed, suppose there exists a solution to the system of equation that satisfies the transversality condition, then we have from the definition of the HJB equation that $\mathrm{E}_t\left[de^{-\rho t}J(x_t) + e^{-\rho t}\{dW_t - \frac{1}{2}\gamma n_t^\top\Sigma n_t dt - \frac{1}{2}J_x^\top\Omega J_x dt - \gamma n_t^\top\Sigma_{sx}J_x dt\}\right] \leq 0 \ \forall n_t$ with equality at the optimal strategy. This implies that $J(x_0) \geq \mathrm{E}[e^{-\rho T}J(x_T)] + \mathrm{E}\left[\int_0^T e^{-\rho t}\{dW_t - \frac{1}{2}\gamma n_t^\top\Sigma n_t dt - \frac{1}{2}J_x^\top\Omega J_x dt - \gamma n_t^\top\Sigma_{sx}J_x dt\}dt\right] \ \forall n$ and with equality for the optimal n. Letting $T \to \infty$ using the bounded convergence theorem and the transversality condition establishes the optimality of the trading strategy and of the value function.

function $J(n,x)$ must satifsy:

$$J(n_t,x_t) = \max_\theta \mathrm{E}_t\left[\int_t^\infty e^{-\rho(s-t)}\left\{dW_u - \frac{1}{2}\gamma n_u^\top \Sigma n_u du - \frac{1}{2}J_x^\top \Omega J_x du - \gamma n_u^\top \Sigma_{sx} J_x du\right\}\right] \tag{149}$$

where we define:

$$\Omega = \gamma\sigma_{xs}\sigma_{xs}^\top + \gamma_x\sigma_x\sigma_x^\top \tag{150}$$

$$\Sigma_{sx} = \sigma_s\sigma_{xs}^\top \tag{151}$$

Thus $J(n,x)$ satisfies the HJB equation:

$$0 = \max_\theta \mathrm{E}_t[dW_t - \frac{1}{2}\gamma n_t^\top \Sigma n_t dt - \frac{1}{2}J_x^\top \Omega J_x dt - \gamma n_t^\top \Sigma_{sx} J_x dt + dJ(n_t,x_t) - \rho J(n_t,x_t)dt] \tag{152}$$

Using the dynamics of the wealth process, we obtain the following equation:

$$0 = \max_\theta\left\{n^\top(\mu_0+\mu x) - \frac{1}{2}\theta^\top\Lambda\theta - \frac{1}{2}\gamma n^\top\Sigma n - \frac{1}{2}J_x^\top\Omega J_x - \gamma n^\top\Sigma_{sx}J_x + J_n^\top\theta - J_x^\top\kappa x + \frac{1}{2}\mathrm{Tr}(J_{xx}\Sigma_x) - \rho J\right\}$$

and we have defined J_x and J_{xx} as respectively the gradient and hessian of $J(n,x,t)$ with respect to x, and J_n the gradiant with respect to n.

The first order condition is

$$\theta = \Lambda^{-1}J_n$$

Plugging back into the HJB equation we get:

$$0 = \max_\theta\left\{n^\top(\mu_0+\mu x) + \frac{1}{2}J_n^\top\Lambda^{-1}J_n - \frac{1}{2}\gamma n^\top\Sigma n - \frac{1}{2}J_x^\top\Omega J_x - \gamma n^\top\Sigma_{sx}J_x - J_x^\top\kappa x + \frac{1}{2}\mathrm{Tr}(J_{xx}\Sigma_x) - \rho J\right\}$$

We guess that the value function is of the form:

$$J(n,x) = -\frac{1}{2}n^\top Q n + n^\top(q_0 + q^\top x) + c_0 + c_1^\top x + \frac{1}{2}x^\top c_2 x$$

where Q, c_2 are symmetric (w.l.o.g.) square matrices and q_0, q_1, c_0, c_1 are all matrices or vectors (with appropriate dimensions) of constant coefficients.

$$\begin{aligned} J_n &= -Qn + q_0 + q^\top x \\ J_x &= qn + c_1 + c_2 x \\ J_{xx} &= c_2 \end{aligned}$$

Thus HJB becomes

$$
\begin{aligned}
0 = &-\rho(\frac{1}{2}n^\top Qn - n^\top(q_0 + q^\top x) - c_0 - c_1{}^\top x - \frac{1}{2}x^\top c_2 x) \\
&+ \frac{1}{2}(-Qn + q_0 + q^\top x)^\top \Lambda^{-1}(-Qn + q_0 + q^\top x) + n^\top(\mu_0 + \mu x) - \frac{1}{2}\gamma n^\top \Sigma n \\
&- \frac{1}{2}(qn + c_1 + c_2^\top x)^\top \Omega(qn + c_1 + c_2 x) - \gamma n^\top \Sigma_{sx}(qn + c_1 + c_2 x) - (qn + c_1 + c_2 x)^\top \kappa x + \frac{1}{2}\operatorname{Tr}(c_2 \Sigma_x)
\end{aligned}
$$

Rewriting:

$$
\begin{aligned}
0 =&\frac{1}{2}n^\top(-\rho Q + Q\Lambda^{-1}Q - \gamma\Sigma - q^\top \Omega q - 2\gamma\Sigma_{sx}q)n \\
&+ n^\top(\rho q_0 + \mu_0 - Q\Lambda^{-1}q_0 - q^\top\Omega c_1 - \gamma\Sigma_{sx}c_1) + n^\top(\rho q^\top - Q\Lambda^{-1}q^\top + \mu - q^\top\kappa - q^\top\Omega c_2 - \gamma\Sigma_{sx}c_2)x \\
&+ x^\top(\frac{1}{2}\rho c_2 + \frac{1}{2}q\Lambda^{-1}q^\top - c_2\kappa - \frac{1}{2}c_2\Omega c_2)x \\
&+ x^\top(\rho c_1 + q\Lambda^{-1}q_0 - c_2\Omega c_1 - \kappa^\top c_1) + \rho c_0 + \frac{1}{2}q_0^\top\Lambda^{-1}q_0 - \frac{1}{2}c_1^\top\Omega c_1 + \frac{1}{2}\operatorname{Tr}(c_2\Sigma_x)
\end{aligned}
$$

So we obtain the set of ODEs that need to be satisfied by the solution.

$$
\rho Q = \gamma\Sigma - Q\Lambda^{-1}Q + q^\top\Omega q + \gamma(\Sigma_{sx}q + q^\top\Sigma_{sx}^\top) \tag{153}
$$

$$
\rho q^\top = \mu - q^\top\kappa - Q\Lambda^{-1}q^\top - q^\top\Omega c_2 - \gamma\Sigma_{sx}c_2 \tag{154}
$$

$$
\rho c_2 = -(c_2\kappa + \kappa^\top c_2) + q\Lambda^{-1}q^\top - c_2\Omega c_2 \tag{155}
$$

$$
\rho c_0 = \frac{1}{2}\operatorname{Tr}(c_2\Sigma_x) + \frac{1}{2}q_0^\top\Lambda^{-1}q_0 - \frac{1}{2}c_1^\top\Omega c_1 \tag{156}
$$

$$
\rho q_0 = \mu_0 - Q\Lambda^{-1}q_0 - q^\top\Omega c_1 - \gamma\Sigma_{sx}c_1 \tag{157}
$$

$$
\rho c_1 = -\kappa^\top c_1 + q\Lambda^{-1}q_0 - c_2\Omega c_1 \tag{158}
$$

We note that if $\mu_0 = 0$ then $c_1 = 0$ and $q_0 = 0$. Also, if $\Omega = 0$ (for example in the GP case, where there is no correlation $\Sigma_{xs} = 0$ and there is vanishing risk-aversion to x risk, that is $\gamma_x = 0$) then the system for Q, q is autonomous and does not depend on the solution for c_2, whereas when there is a hedging demand $\gamma_x > 0$ then the system for Q, q, c_2 needs to be solved jointly. So c_2 encodes the hedging demand component, just like in the case without transaction costs.

To interpret the optimal trading strategy, note that the value function is maximized with respect to the position vector n at the optimal aim portfolio $aim(x_t) = Q^{-1}(q_0 + q^\top x_t)$. Since $J_n = -Qn + q_0 + q^\top x$ the optimal trade can be written as:

$$
\theta = \Lambda^{-1}J_n = \Lambda^{-1}Q(aim(x_t) - n_t)
$$

So with the definition of fixed trade intensity $\tau = \Lambda^{-1}Q$ we get the optimal trading strategy:

$$dn_t = \tau(aim(x_t) - n_t)dt \tag{159}$$

J The one asset one factor case

Here we analyze the solution for the simple special case of one-factor and one asset, that is $N = K = 1$. We further set $\mu_0 = 0$.

J.1 The infinite horizon no-transaction-cost case

We note that the if $\mu_0 = 0$ then $c_1 = 0$ and the trading strategy only depends on c_2 which solves the quadratic equation:

$$0 = c_2 \left(\gamma\Sigma_{sx}^\top\Sigma^{-1}\Sigma_{sx} - \Omega\right) c_2 - c_2(\rho + 2\kappa + 2\Sigma_{sx}\Sigma^{-1}\mu) + \mu^\top(\gamma\Sigma)^{-1}\mu \tag{160}$$

Recall that $\Omega = \gamma\sigma_{xs}^2 + \gamma_x\sigma_x^2$ and $\Sigma_{sx} = \sigma_s\sigma_{xs}$ and $\Sigma = \sigma_s^2$. Thus the equation simplifies:

$$0 = c_2^2\gamma_x\sigma_x^2 - c_2(\rho + 2\kappa + 2\frac{\sigma_{xs}}{\sigma_s}\mu) + \frac{\mu^2}{\gamma\sigma_s^2} \tag{161}$$

The solution is given in the main paper.

J.2 The infinite horizon with tcost

We note that the if $\mu_0 = 0$ then $c_1 = q_0 = 0$ and the trading strategy only depends on c_2, Q, q which solve the equations:

$$\rho Q = \gamma\Sigma - Q\Lambda^{-1}Q + q^\top\Omega q + 2\gamma\Sigma_{sx}q \tag{162}$$

$$\rho q^\top = \mu - q^\top\kappa - Q\Lambda^{-1}q^\top - q^\top\Omega c_2 - \gamma\Sigma_{sx}c_2 \tag{163}$$

$$\rho c_2 = -2c_2\kappa + q\Lambda^{-1}q^\top - c_2\Omega c_2 \tag{164}$$

Recall that $\Omega = \gamma\sigma_{xs}^2 + \gamma_x\sigma_x^2$ and $\Sigma_{sx} = \sigma_s\sigma_{xs}$ and $\Sigma = \sigma_s^2$. Thus the equations simplify:

$$\rho Q = \gamma\Sigma - Q\Lambda^{-1}Q + q^2\Omega + 2\gamma\Sigma_{sx}q \tag{165}$$

$$0 = \mu - q(\kappa + \rho) - Q\Lambda^{-1}q - q\Omega c_2 - \gamma\Sigma_{sx}c_2 \tag{166}$$

$$0 = -(\rho + 2\kappa)c_2 + q^2\Lambda^{-1} - c_2^2\Omega \tag{167}$$

We now express everything in terms of the trading speed: $\tau = \Lambda^{-1}Q$ to get:

$$\tau^2 + \rho\tau = \gamma\Lambda^{-1}\Sigma + q^2\Lambda^{-1}\Omega + 2\gamma\Lambda^{-1}\Sigma_{sx}q \tag{168}$$

$$0 = \mu - q(\kappa + \rho) - \tau q - q\Omega c_2 - \gamma\Sigma_{sx}c_2 \tag{169}$$

$$0 = (\rho + 2\kappa)c_2 - q^2\Lambda^{-1} + c_2^2\Omega \tag{170}$$

To solve this problem, we see that the first equation has a unique positive root for $\tau(q)$ and the last equation has a unique positive root $c_2(q)$, both given by:

$$\tau(q) = \frac{-\rho + \sqrt{\rho^2 + 4\Lambda^{-1}(q^2\gamma_x\sigma_x^2 + \gamma(\sigma_s + q\sigma_{xs})^2)}}{2} \tag{171}$$

$$c_2(q) = \frac{-(\rho + 2\kappa) + \sqrt{(\rho + 2\kappa)^2 + 4\Lambda^{-1}q^2(\gamma\sigma_{xs}^2 + \gamma_x\sigma_x^2)}}{2\Omega} \tag{172}$$

The solution is then found by solving the second equation for q.

$$q(\kappa + \rho) + \tau(q)q + (q\gamma_x\sigma_x^2 + \gamma\sigma_{xs}(q\sigma_{xs} + \sigma_s))c_2(q) = \mu \tag{173}$$

It is clear that this equation always admits a positive solution (since the left hand side is a continuous function equal to zero when $q = 0$ and tending to infinity as $q \to \infty$).

The optimal aim portfolio is given by

$$aim(x) = Q^{-1}qx \tag{174}$$

$$= \tau^{-1}\Lambda^{-1}q \tag{175}$$

$$\tag{176}$$

K The infinite horizon solution CMV preferences (i.e., $\sigma_{xs} = 0$ and $\gamma_x = 0$)

As discussed in Remark 2, the solution to the finite horizon model where agents have CMV preferences (as in equation (34)) corresponds to the solution of the source-dependent risk-aversion recursive utility agent with parameters restricted to $\sigma_{xs} = 0$ and $\gamma_x = 0$ (which implies $\Omega = 0$).

To understand the optimal trading rule (aim, τ) the relevant system of ODE we need to solve becomes:

$$\rho Q = \gamma\Sigma - Q\Lambda^{-1}Q \tag{177}$$

$$\rho q^\top = \mu - q^\top\kappa - Q\Lambda^{-1}q^\top \tag{178}$$

$$\rho q_0 = \mu_0 - Q\Lambda^{-1}q_0 \tag{179}$$

Now, we can rewrite this system in terms of the trading speed matrix $\tau = \Lambda^{-1}Q$ as:

$$\rho\tau = \gamma\Lambda^{-1}\Sigma - \tau\tau \tag{180}$$

$$\rho Q = \gamma\Sigma - \tau^\top Q \tag{181}$$

$$\rho q^\top = \mu - q^\top\kappa - \tau^\top q^\top \tag{182}$$

$$\rho q_0 = \mu_0 - \tau^\top q_0 \tag{183}$$

This system has an intuitive closed-form solution. Let us define the diagonalization $\gamma\Lambda^{-1}\Sigma = FD_\eta F^\top$, where we define D_η as the diagonal matrix with eigenvalue η_i on the i^{th} diagonal. Then we see that $\tau_t = FD_hF^\top$ where h_i are the positive root of the quadratic equation:

$$\rho h_i = \eta_i - h_i^2 \tag{184}$$

The solution is:

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

It follows that the trading speed matrix is given by $\tau = FD_hF^\top$ and the the Q matrix is $Q = \Lambda\tau$.

To solve for q, q_0, we note that they can be expressed directly in terms of the trading speed and using lemma 3 for the expression for q

$$Q = (\rho + \tau^\top)^{-1}\gamma\Sigma \tag{185}$$

$$q_0 = (\rho + \tau^\top)^{-1}\mu_0 \tag{186}$$

$$q^\top = \int_0^\infty e^{-(\rho+\tau^\top)t}\mu e^{-\kappa t}dt \tag{187}$$

lemma 3 *Suppose A, B are (full rank) square matrices with strictly positive eigenvalues. Then the matrix equation $-AX - XB = -C$ has the solution $X = \int_0^\infty e^{-At}Ce^{-Bt}dt$.*

Proof. Note that:

$$e^{-AT}Ce^{-BT} - C = \int_0^T d(e^{-At}Ce^{-Bt}) = -A\int_0^T e^{-At}Ce^{-Bt}dt - \int_0^T e^{-At}Ce^{-Bt}dtB$$

Taking the limit as $T \to \infty$ and noting that, since all the eigenvalues of A, B are positive we have $\lim_{T\to\infty} e^{-AT}Ce^{-BT} = 0$, we obtain: $-C = -AX - XC$ where X is as defined in the lemma. ■

The main result then follows from the definition of the aim portfolio $aim(x) = Q^{-1}(q_0 + qx)$.