

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

Pierre Collin-Dufresne
EPFL and SFI and CEPR
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: September 2024

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 future 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. We illustrate the differences in portfolio makeup that result from considering hedging demands of a long-horizon investor using a set of simplified examples, and using a daily trading strategy based on the estimated relation between retail order imbalance and future returns.

JEL Classification: D53, G11, G12

Keywords: portfolio choice, dynamic models, transaction costs, hedging demand, price impact, mean-variance

*We thank Agostino Capponi, Darrell Duffie, Julien Hugonnier, Kevin Webster and the participants of the Workshop on Non-Standard Investment Choice at ESSEC, the Princeton University Stochastic Control and Financial Engineering Conference, the Peking University Math Finance Seminar, and the 2002 Financial Management Association Meetings for helpful discussions and comments. 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, this portfolio hedges the investor against changes in the investment opportunity set.

Second, if there are transaction costs then it will not be optimal to continuously and fully rebalance a portfolio in response to 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

While the CMV objective function has the advantage of tractability, it has the peculiar implication that agents with CMV preferences do not care about the correlation between stock returns and the investment opportunity set and therefore display no hedging demands as defined in Merton (1969, 1971) and Cox and Huang (1989). Capturing such non-myopic behavior requires longer-term risk aversion, for example with a preference specification in which investors maximize their

¹Collin-Dufresne, Daniel, and Saġlam (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.

²In 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).

expected utility over their terminal wealth.

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), the agent exhibits differential risk aversion to shocks to prices and shocks to expected returns. We show that this preference specification is equally as tractable as the CMV framework.

The preference specification facilitates comparison of the optimal dynamic portfolio choices that arise from the different, nested, preference specifications. Specifically, in the finite horizon case, when risk-aversion coefficients towards all sources of risk are equal, this preference specification nests standard CARA (negative exponential) expected utility.

Interestingly, the stationary CMV specification used in the literature, summarized by the objective (\star) above, also corresponds to maximization of certainty equivalent wealth for a source-dependent recursive-utility investor in a setting where the horizon at which final wealth is evaluated is drawn from an exponential distribution with parameter ρ , but in the limit where the agent approaches risk-neutrality toward the risks driving expected returns and toward horizon risk.

We characterize the closed-form solution to the optimal portfolio choice problem with this generalized preference specification 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. Doing so allows us to characterize how transactions costs and hedging demands will affect an investor’s (optimal) trading decisions.

We show that, even for the general preferences considered, the agent’s optimal policy is to trade towards an aim portfolio at a given trading speed.

As in GP, we find that for the agent with (myopic) CMV preferences the aim portfolio is a trading-speed-discounted average of expected future CMVE portfolios where the optimal trading speed matrix is entirely determined by the ratio of the stock volatility matrix to the price impact matrix. Thus a CMV-investor will always optimally underweight a security with a mean-reverting expected return relative to their weight in the CMVE portfolio, where by “underweight” we mean that the weight is closer to zero.³ Further, the CMV-investor’s aim portfolio and trading speed are independent of the correlation structure of signals. Indeed, a CMV-investor makes identical portfolio choices whether signals are deterministic or stochastic.

Instead, when the agent is a long-term (CARA) expected utility investor, the ability to hedge changes in the investment opportunity set can dramatically affect both the composition of the aim portfolio and the trading speed. Both depend crucially on the correlation between realized stock returns and shocks to expected returns. If this correlation is negative, then a long term CARA investor will typically choose to overweight a stock relative to the CMVE benchmark, despite it having high transaction costs and mean-reverting return. In fact, if the correlation is sufficiently negative, we find that the long-term investor might even choose to overweight an asset relative to her no-tcost optimal benchmark (which itself is typically overweight relative to the CMVE due to

³This weight is, however, *entirely independent* of the covariance matrix of signals.

the classic Merton hedging demand).

Furthermore, the speed at which the agent will optimally trade towards the aim portfolio depends on the correlation between shocks to expected returns and realized stock returns. In particular, the trading speed is lower when realized stock returns are more negatively correlated with shocks to expected returns. 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. Since transaction costs increase the long-term investor’s cost of hedging, this decreases the optimal trading speed. Finally, since the aim portfolio is a trading-cost-discounted average of future expected no-transaction-cost portfolios, a negative correlation also implies a smaller discount between the no-transaction-cost optimum and the aim portfolio.

As an empirical application, we calibrate a model based on the findings of Boehmer, Jones, Zhang, and Zhang (2021), who develop an algorithm to identify retail trades and find empirical evidence that stocks with net retail buying outperform those with net selling over the following week. In this setting, we explore the utility benefits of employing our approach with hedging demands when compared to a CMV investor who fundamentally ignores the correlation between the innovations in asset returns and predictors. We focus our analysis on the 25 largest U.S. stocks from 2014 to 2019, finding that daily retail order flow positively predicts next-day returns for several stocks with statistical significance. We observe that the predictive signal exhibits rapid mean reversion, with an average half-life of 1.2 days, and there exists a significant negative correlation between return innovations and predictor innovations, consistent with contrarian retail trading behavior. Finally, we calibrate the transaction cost model using a large institutional large order data set from the global execution desk of a large investment bank.

Using this realistically calibrated model, we compare the performance of CARA and CMV strategies across various experiments, varying the number of stocks (2 to 25), investment horizons (3 to 6 months), and risk aversion levels. Our results consistently show that the CARA strategy that incorporates hedging demands, outperforms the CMV strategy that ignores the correlation between innovations in asset returns and predictors. We find these utility gains to be both statistically and economically significant, with improvements in certainty equivalent wealth of up to 7% over the examined horizons. We find that the CARA strategy generally takes larger positions in stocks that exhibit a stronger negative correlation between return and predictor innovations, higher signal predictability (as indicated by its Sharpe ratio), and lower mean-reversion in the predictor dynamics. Overall, our findings demonstrate substantial economic value in accounting for hedging demands even in a problem with short-term investment horizon.

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 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.⁴ Finally, Muhle-Karbe, Sefton, and Shi (2023) extend the GP framework by adjusting its objective function to maximize the lifetime utility of consumption for an agent with CARA preferences. Their primary focus is on the existence of a solution to the coupled Riccati equations, and they provide a rigorous verification theorem that correctly identifies the value function, along with the optimal consumption and trading policy. Interestingly, they show that the resulting trading speed arises from the solution of an optimal execution problem (Almgren and Chriss, 2001).

⁴It would be interesting to extend our model to study how hedging demands driven by stochastic shifts in second moments affect their findings.

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, for simplicity, that the risk-free rate is zero.⁵ Z^s and Z^x are vectors of independent Brownian motions that drive the randomness in stock prices and the state variables.⁶ 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:*

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

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

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 .⁷

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:

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

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

⁵For 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. The zero risk-free rate assumption could easily be relaxed to constant or affine in x .

⁶Since dZ_t^s is $N \times 1$, σ_x is $K \times K$, σ_s is $N \times N$ and σ_{xs} is $K \times N$.

⁷Other matrices need to be adjusted appropriately as well. For example, μ is the (N, NM) diagonal sparse matrix which has row vector $[\mu_{1,i}, \mu_{2,i}, \dots, \mu_{N,i}]$ on the i^{th} ‘diagonal.’

where Λ is a symmetric positive definite transaction-cost matrix.⁸

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):

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

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 (expected future) 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:

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

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

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

Proof. See Appendix A and Appendix B. ■

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 shows, 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 *For general x_t process, the CMV objective function of equation (9) reduces to the linear-quadratic framework used in Litterman (2005), Gârleanu and Pedersen (2013), and Collin-*

⁸Assuming 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.

Dufresne, Daniel, and Sağlam (2020):

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

The optimal value-function of a CMV-investor 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 $\sigma_{xs} = 0$. Further, the optimal portfolio choice of a CMV-investor who maximizes (9) or equivalently (10) for arbitrary σ_x and σ_{xs} is independent of σ_x and σ_{xs} and thus identical to that of a source-dependent utility agent who maximizes (7) with $\gamma_x \sigma_x = 0$ and $\sigma_{xs} = 0$. In other words, **the CMV-agent acts as if expected returns were deterministic.**

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).⁹

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.

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

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

⁹See Theorem 5 below for a precise restatement of this result in the context of our model.

where the vector $c_1(t)$ and the symmetric matrix $c_2(t)$ solve the system of ODEs:

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

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

where

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

$$\Sigma_{sx} = \sigma_s \sigma_{xs}^\top, \quad (15)$$

and the boundary conditions are given by $c_1(T) = 0$ and $c_2(T) = 0$. In particular, if $\mu_0 = 0$ then $c_1(t) = 0 \forall t$. $c_0(t)$ is given in equation (97) in Appendix E.

The optimal position (in the absence of transaction costs) is given by:

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

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

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

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

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

where $Q(t)$ and $c_2(t)$ are symmetric (respectively N - and K -dimensional) matrices, $q(t)$ is a $(K \times N)$ matrix, $q_0(t)$ and $c_1(t)$ are vectors that solve the system of ODEs:

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

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

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

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

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

subject to boundary conditions $Q(T) = 0$, $q(T) = 0$, $q_0(T) = 0$, $c_1(T) = 0$ and $c_2(T) = 0$. $c_0(t)$ is given in equation (111) 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:

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

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

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

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

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 $\text{aim}(x_t, t)$ at a rate of τ_t .

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 obtained in GP. In our continuous time setting, it can be fully characterized in terms of the eigenvalue decomposition of the matrix $\gamma\Lambda^{-1}\Sigma$. Specifically, we assume the following:

Assumption 1 *The square matrix $\gamma\Lambda^{-1}\Sigma$ has full rank and distinct real eigenvalues, so it can be diagonalized:*

$$\gamma\Lambda^{-1}\Sigma = FD_\eta F^{-1} \quad (27)$$

where D_η is the $N \times N$ diagonal matrix with eigenvalue η_i on the i^{th} diagonal, and F is the corresponding square 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:*

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

The optimal aim portfolio of the investor with CMV preferences given in equation (26) can be interpreted as a Markowitz portfolio where we replace the expected return vector by a trading-speed

weighted average of future expected returns:

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

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

where the expected future stock return is defined by

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

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,¹⁰ 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).

Below we will compare our general solution obtained for the non-myopic CARA agent in theorem 4 with that of the CMV-agent in theorem 5 using specific cases, numerical examples, and one specific empirical implementation, to illustrate the importance of hedging demands on portfolio choices in the context of transaction costs.

But first, 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. 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. 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.

¹⁰Recall that $E_t[x_u] = e^{-\int_t^u \kappa ds} x_t$.

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):

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

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

One might think that this stationary version of equation (7) should correspond to the certainty equivalent of a CARA agent who maximizes $\mathbb{E}[-e^{-\gamma W_{\mathcal{T}}}]$ 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 $(Z^s, Z^x, \mathbf{1}_{\{\mathcal{T} \leq t\}})$, consider the process $(H_t, \sigma_{H,s}, \sigma_{H,x})$ which solves the following backward stochastic differential equation (BSDE):*

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

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, $\mathbf{1}_{\{\mathcal{T} \leq t\}}$, 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:

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

- 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 $\sigma_{x_s} = 0$, it reduces to the discounted CMV objective function:

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

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*

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

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

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

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:

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

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 symmetric positive definite matrix c_2 and the vector c_1 solve the system of quadratic

equations:

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

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

In particular, if $\mu_0 = 0$, then $c_1 = 0$. The equation for c_0 is provided in the Appendix.

The optimal position (in the absence of transaction costs) is:

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

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

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

where Q and c_2 are positive-definite symmetric (respectively $N \times N$ and $K \times K$) matrices, q is a $N \times K$ matrix, and q_0, c_1 are (respectively N - and K -dimensional) vectors that solve the following system of quadratic equations:

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

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

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

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

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

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:

$$dn_t = \tau (\text{aim}(x_t) - n_t) dt \quad (45)$$

$$\tau = \Lambda^{-1} Q \quad (46)$$

$$\text{aim}(x) = Q^{-1} (q_0 + q^\top x) \quad (47)$$

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

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 the absence of transaction costs, the CMV investor always holds the CMVE portfolio. With transaction costs however, the solution of the CMV investor can be characterized explicitly (setting $\gamma_x = 0$ and $\sigma_{xs} = 0$ in theorem 8), in terms of the eigenvalue and eigenvector decomposition (η, F) of the matrix $\gamma\Lambda^{-1}\Sigma$.

Indeed, we have the following result, that is the infinite horizon stationary equivalent to theorem 5:

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

$$\tau = FD_h F^{-1}$$

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

The optimal aim portfolio of the GP investor of equation (47) can be written as a Markowitz portfolio where we replace the instantaneous expected stock return $\mu_S(x_t) = \frac{1}{dt}\mathbf{E}_t[dS_t] = \mu_0 + \mu x_t$ by the trading speed discounted value of the future stock expected returns:

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

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

Proof. The proof is in 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 CMVE-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:

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

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

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

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-CMVE portfolio and the CMV aim portfolio. Further, 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 $(\gamma\Sigma)^{-1}(\mu_0 + \mu x_t) - \Sigma^{-1}\Sigma_{sx}(c_1 + c_2x)$. In the presence of transaction costs and with $\sigma_{xs} \neq 0$, the equation is more difficult to interpret, especially in the multi-asset case, because the covariance matrix of signals and returns affects trading speed and aim portfolios. While the equation has a similar structure, which suggests that the intuition of discounting future no-tcost optimal portfolios, that themselves contain a hedging demand, remains useful the actual results are more complex (in particular, because 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:

$$dS(t) = \mu x_t dt + \sigma_1 dZ_1(t) \quad (48)$$

$$dx_t = -\kappa x(t)dt + \sigma_{x1}dZ_1(t) + \sigma_{x2}dZ_2(t) \quad (49)$$

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

We can first solve for the optimal portfolio of the non-myopic agent in the stationary case using Theorem 8. If transaction costs are zero (if $\Lambda = 0$), the optimal portfolio is a combination of the *CMVE* portfolio and a hedging portfolio *HP*:

$$n_t = CMVE_t + HP_t \quad (50)$$

$$CMVE_t = \frac{\mu}{\gamma\sigma_1^2} x_t \quad (51)$$

$$HP_t = -\frac{2\left(\frac{\mu}{\sigma_1}\right)^2 \frac{\sigma_{x1}}{\sigma_1}}{\gamma \left(2\kappa + \rho + 2\frac{\mu}{\sigma_1}\sigma_{x1} + \sqrt{(2\kappa + \rho + 2\frac{\mu}{\sigma_1}\sigma_{x1})^2 + 4\frac{\gamma_x}{\gamma} \left(\frac{\mu}{\sigma_1}\right)^2 \sigma_{x2}^2}\right)} x_t \quad (52)$$

As expected the non-myopic agent deviates from the CMVE portfolio if and only if innovations in expected returns are correlated with the realized returns of the risky asset: that is if $\sigma_{x1} \neq 0$. If $\sigma_{x1} < 0$, then falls in the price of the risky asset generally lead to an increase in future expected

returns, making it less risky from a long-term perspective. This will lead the agent to scale up her investment in the risky asset.

In a setting where the agent has myopic CMV preferences, the agent will deviate from the CMVE portfolio only if t-costs are positive. Applying Theorem 9 we can derive the optimal aim portfolio and trading speed as follows:

$$aim_t^{CMV} = \left(\frac{\rho + \tau}{\rho + \tau + \kappa} \right) \frac{\mu}{\gamma \sigma_1^2} x_t \quad (53)$$

$$\tau = \frac{1}{2} \left(\sqrt{\rho^2 + 4\gamma \frac{\sigma_1^2}{\lambda^2}} - \rho \right) \quad (54)$$

The CMV-agent's aim portfolio is the CMVE portfolio only if $\kappa = 0$, otherwise her holdings are strictly 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:

$$aim_t^{CARA} = \frac{\rho + \tau}{\rho + \tau + \kappa} \frac{\mu - (\gamma \sigma_1 \sigma_{x1} + \Omega q) c_2}{\gamma \sigma_1^2 + q^2 \Omega + 2\gamma \sigma_1 \sigma_{x1} q} x_t \quad (55)$$

$$\tau = \frac{1}{2} \left(\sqrt{\rho^2 + 4 \left\{ \gamma \left(\frac{\sigma_1}{\lambda} + \frac{q}{\lambda} \sigma_{x1} \right)^2 + \frac{q^2}{\lambda^2} \gamma_x \sigma_{x2}^2 \right\}} - \rho \right) \quad (56)$$

$$c_2 = \frac{\sqrt{(2\kappa + \rho)^2 + 4\Omega \frac{q^2}{\lambda^2}} - \rho - 2\kappa}{2\Omega} \quad (57)$$

where q is the constant that solves the following non-linear equation:¹¹

$$c_2 (\gamma \sigma_1 \sigma_{x1} + \Omega q) + q(\rho + \kappa + \tau) = \mu \quad (58)$$

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

¹¹Note 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 $q \rightarrow \infty$ (see Appendix J for details).

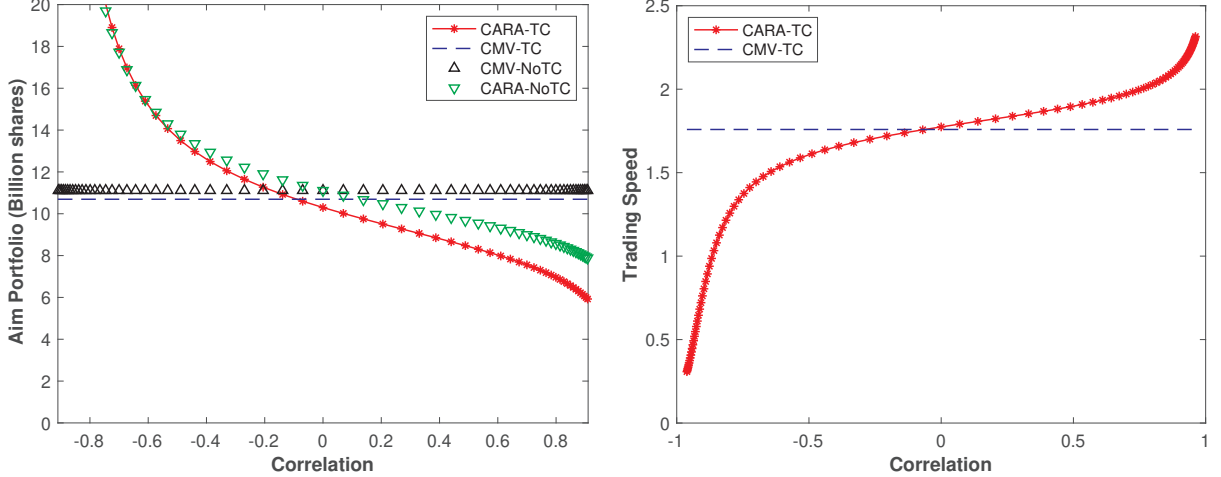


Figure 1: Aim portfolios and trading speeds in low transaction cost regime.

Notes: In the left panel, we plot the aim portfolios for CMV and CARA agents in the presence and absence of transaction costs as a function of the correlation between the innovations in stock and predictor dynamics. In the right panel, we plot the trading speeds as a function of the same correlation for CMV and CARA agents in the presence of transaction costs. Here, the correlation is given by $\frac{\sigma_{x1}}{\sqrt{\sigma_{x1}^2 + \sigma_{x2}^2}}$ and we vary σ_{x1} while keeping σ_{x2} constant. Parameters: $\mu_0 = 0$, $\mu = 1$, $\kappa = 0.1$, $\sigma_1 = 0.3$, $\sigma_{x2} = 0.1$, $\Lambda = 2 \times 10^{-11}$, $\gamma_x = 10^{-9}$, $\gamma = 10^{-9}$, $x_0 = 1$, $\rho = 0.8$.

a non-myopic investor shapes her optimal trading strategy in the presence of transaction costs (CARA-TC). Thus we report the trading strategy of a CMV investor (CMV-TC) 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 (CMV-NoTC), and compare it with that of a source dependent risk-aversion investor with $\gamma_x = \gamma$ who maximizes (32). That is, we examine a CARA investor who is necessarily risk-averse with respect to changes in the investment opportunity set.¹² We also include the aim portfolio for a CARA agent in the absence of transaction costs, labeled as CARA-NoTC.

Figure 1 reports the aim portfolio holdings of the risky-asset in the left panel and the trading speed in the right panel as a function of the correlation between the innovations in stock and predictor dynamics. Here the trading cost is low ($\Lambda = 2 \times 10^{-11}$) and the expected return is positive ($\mu x_0 = 1$). As expected, it shows that the CMV-TC investor's optimal aim portfolio is very close to the mean-variance efficient Markowitz portfolio, CMV-NoTC. 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-TC investor chooses a portfolio very similar to that of the classic no-transaction-cost *Merton* solution, CARA-NoTC. Specifically, she displays a very large and positive hedging demand for the asset when correlation between x and dS becomes negative. This is 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

¹²Note that since we assume $\gamma_T = 0$, the investor we consider is risk-neutral with respect to horizon realization risk.

expected returns and stock price changes. With negative correlation, expected returns changes offer a natural hedge for shocks to stock prices.

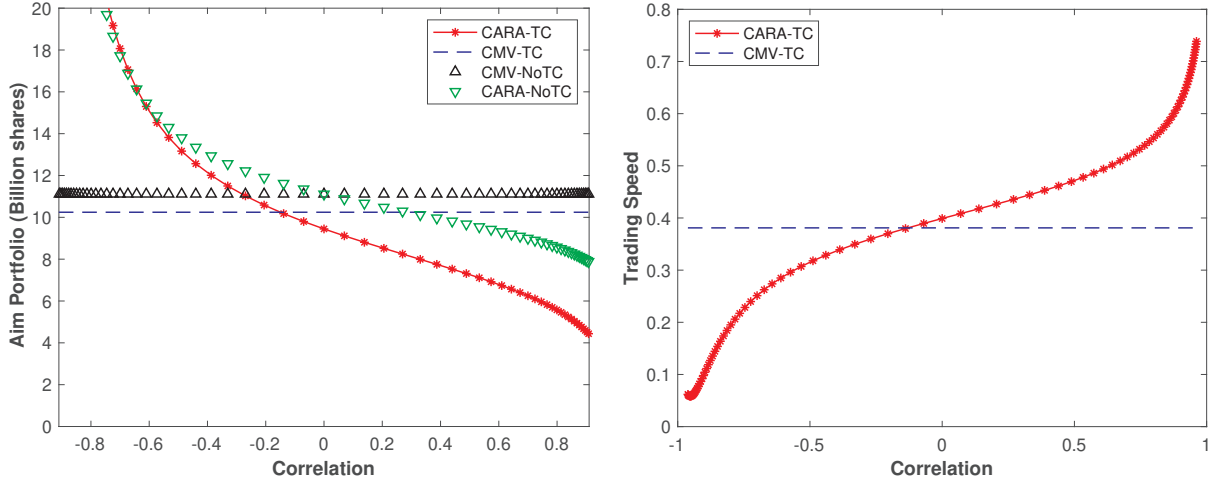


Figure 2: Aim portfolios and trading speeds in high transaction cost regime.

Notes: In the left panel, we plot the aim portfolios for CMV and CARA agents in the presence and absence of transaction costs as a function of the correlation between the innovations in stock and predictor dynamics. In the right panel, we plot the trading speeds as a function of the same correlation for CMV and CARA agents in the presence of transaction costs. Here, the correlation is given by $\frac{\sigma_{x1}}{\sqrt{\sigma_{x1}^2 + \sigma_{x2}^2}}$ and we vary σ_{x1} while keeping σ_{x2} constant. We use the following parameters: $\mu_0 = 0$, $\mu = 1$, $\kappa = 0.1$, $\sigma_1 = 0.3$, $\sigma_{x2} = 0.1$, $\mathbf{\Lambda} = \mathbf{2} \times \mathbf{10}^{-10}$, $\gamma_x = 10^{-9}$, $\gamma = 10^{-9}$, $x_0 = 1$, $\rho = 0.8$.

In Figure 2, all parameters are the same except the trading cost, which is now an order of magnitude higher. In response to the higher t-costs, the CMV-TC investor now chooses an aim portfolio that is uniformly lower than the Markowitz portfolio, CMV-NoTC. Intuitively, because of transaction costs the investor has to trade slowly into 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-TC 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 that of CARA-NoTC occurs for a correlation coefficient (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 surprising 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 paid following the negative stock price shock will be weighted more. Therefore it can be optimal to aim for a higher stock position and trade less to avoid paying the transaction costs in the high marginal utility states (following a negative stock price shock). Note that this intuition suggests that, for a CARA agent, the weight on the risky asset will be scaled down relative to the no-t-cost weight still further when the correlation is positive. This is also observed in the Figure.

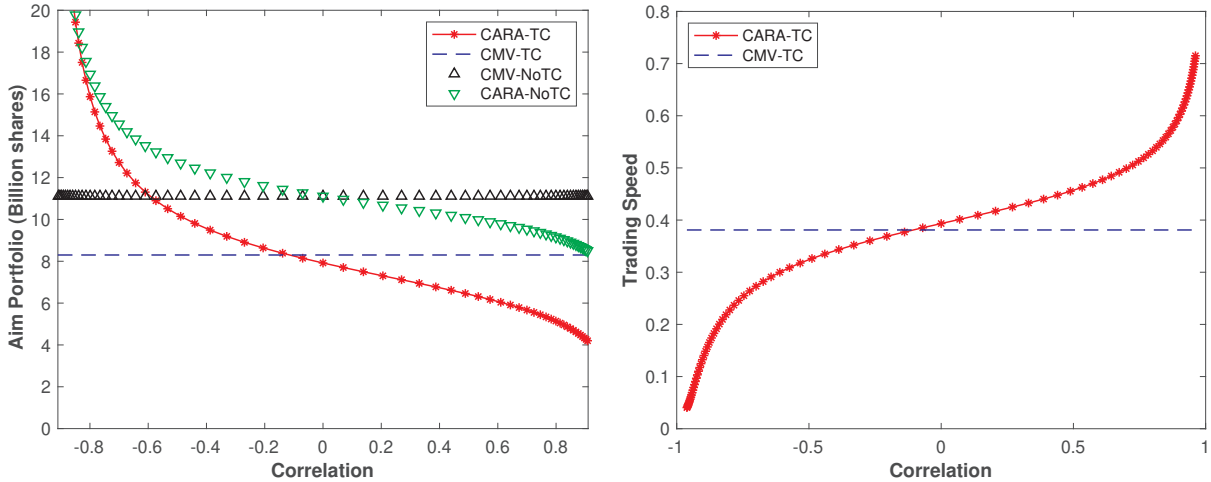


Figure 3: Aim portfolios and trading speeds in high mean-reversion regime for the predictor.

Notes: In the left panel, we plot the aim portfolios for CMV and CARA agents in the presence and absence of transaction costs as a function of the correlation between the innovations in stock and predictor dynamics. In the right panel, we plot the trading speeds as a function of the same correlation for CMV and CARA agents in the presence of transaction costs. Here, the correlation is given by $\frac{\sigma_{x1}}{\sqrt{\sigma_{x1}^2 + \sigma_{x2}^2}}$ and we vary σ_{x1} while keeping σ_{x2} constant. We use the following parameters: $\mu_0 = 0$, $\mu = 1$, $\kappa = 0.4$, $\sigma_1 = 0.3$, $\sigma_{x2} = 0.1$, $\Lambda = 2 \times 10^{-10}$, $\gamma_x = 10^{-9}$, $\gamma = 10^{-9}$, $x_0 = 1$, $\rho = 0.8$.

In Figure 3 we show the effect of having a less persistent signal. Here $\kappa = 0.4$, and all other parameters are the same as in Figure 2. The expected return and variance are the same as in Figures 1 and 2, so the Markowitz CMV-NoTC portfolio is unchanged, but since the expected return will now decay faster the expected return earned over the life of the position will be lower, so the aim portfolio of the CMV-TC investor is scaled down still further. Similarly, the CARA-TC investor in the presence of transaction costs reduces her position relative to the CARA-NoTC investor. Again, though, when the correlation between the shocks to price and expected return

become sufficiently negative (close to -80% in this case), we see that because of hedging demands the CARA-TC investor's aim portfolio becomes larger than what she would choose in the absence of transaction costs. So even for fast decaying parameters, when t-costs are large 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 further illustrate the role of hedging demands in shaping the optimal portfolio choice we consider a very specific setup with two stocks, where one stock will optimally only be held by the non-myopic agent. Specifically, we consider dynamics:

$$dS_1(t) = \mu x_t dt + \sigma_1 dZ_1(t) \quad (59)$$

$$dS_2(t) = \sigma_2 dZ_2(t) \quad (60)$$

$$dx_t = -\kappa x(t)dt + \sigma_{x2} dZ_2(t) \quad (61)$$

where the $Z_i(t)$ are independent Brownian motions. 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:

$$n_t = CMVE_t + HP_t \quad (62)$$

$$CMVE_t = \left[\frac{\mu}{\gamma \sigma_1^2}; 0 \right]^\top x_t \quad (63)$$

$$HP_t = \left[0; -\frac{\mu^2 \sigma_{x2}}{\gamma \sigma_1^2 \sigma_2 (\rho + 2\kappa)} \right]^\top x_t \quad (64)$$

We see that the myopic agent only trades asset 1, but has no demand for asset 2. Because asset 2 has zero expected (excess) return and positive variance its weight in the CMVE portfolio is always zero. However, because asset 2 realized returns are correlated with shocks to asset 1's expected return, a non-myopic agent will hold asset 2 as a hedge against changes in asset 1's expected return. 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's realized returns have zero correlation with shocks to its expected return, it is not useful as a hedge.

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 agent. From Theorem 9, the CMV agent's aim portfolio and trading speed are:

$$aim_t^{CMV} = \left[\frac{\rho + \tau_{11}}{\rho + \tau_{11} + \kappa \gamma \sigma_1^2} \frac{\mu}{\lambda_1^2}; 0 \right]^\top x_t \quad (65)$$

$$\tau_{11} = \frac{1}{2} \left(\sqrt{\rho^2 + 4\gamma \frac{\sigma_1^2}{\lambda_1^2}} - \rho \right) \quad (66)$$

$$\tau_{22} = \frac{1}{2} \left(\sqrt{\rho^2 + 4\gamma \frac{\sigma_2^2}{\lambda_2^2}} - \rho \right) \quad (67)$$

$$\tau_{12} = \tau_{21} = 0 \quad (68)$$

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 matrix is diagonal, meaning that the weight on asset 2 do not affect trading of 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.

Even when t-costs are positive, the CMV agent will again have a zero weight on asset 2 in their aim portfolio; if endowed with a position, she would trade out of it. 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 a (non-myopic) CARA agent. Since a CARA agent has a long investment horizon, the covariance of expected return shocks and realized returns will affect both the makeup of her aim portfolio and the speed at which she will trade towards this portfolio.

Solving the system for the optimal aim portfolio and trading speeds results in a system of non-linear equations:

$$aim_t^{CARA} = Q^{-1} q x_t \quad (69)$$

$$Q = \Lambda \tau \quad (70)$$

$$\tau_{11} = \frac{1}{2} \left(\sqrt{\rho^2 + 4\gamma \frac{\sigma_1^2}{\lambda_1^2} + 4\gamma \frac{q_1^2}{\lambda_1^2} \sigma_{x_2}^2 - 4 \frac{\lambda_1^2}{\lambda_2^2} \tau_{12}^2} - \rho \right) \quad (71)$$

$$\tau_{22} = \frac{1}{2} \left(\sqrt{\rho^2 + 4\gamma \left(\frac{\sigma_2}{\lambda_2} + \frac{q_2}{\lambda_2} \sigma_{x_2} \right)^2 - 4 \frac{\lambda_1^2}{\lambda_2^2} \tau_{12}^2} - \rho \right) \quad (72)$$

$$\tau_{12} = \frac{\lambda_2^2}{\lambda_1^2} \tau_{21} = \frac{\lambda_2}{\lambda_1} \frac{\gamma \frac{\sigma_2 q_1}{\lambda_2 \lambda_1} \sigma_{x_2} + \gamma \frac{q_1 q_2}{\lambda_1 \lambda_2} \sigma_{x_2}^2}{\rho + \tau_{11} + \tau_{22}} \quad (73)$$

and where the q_1, q_2, c_2 solve

$$c_2 = \frac{-\rho - 2\kappa + \sqrt{(\rho + 2\kappa)^2 + 4\left(\frac{q_1^2}{\lambda_1^2} + \frac{q_2^2}{\lambda_2^2}\right)\gamma\sigma_{x2}^2}}{2\gamma\sigma_{x2}^2} \quad (74)$$

$$\mu = q_1(\kappa + \rho + c_2\gamma\sigma_{x2}^2 + \tau_{11}) + q_2\frac{\lambda_1^2}{\lambda_2^2}\tau_{12} \quad (75)$$

$$-c_2\sigma_2\gamma\sigma_{x2} = q_2(\kappa + \rho + c_2\gamma\sigma_{x2}^2 + \tau_{22}) + q_1\tau_{12} \quad (76)$$

We solve this system numerically and show how aim portfolio changes with the parameters 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 examined earlier. Figure 4 plots the aim portfolio holdings of asset 1 (left panel) and asset 2 (right panel) as a function of σ_{x2} .

The aim portfolio weight on asset 1 is not affected by t-costs for either the CMV-TC or the CARA-TC agent. The weight on asset 2 is zero for the myopic CMV-NoTC agent, but the non-myopic CARA-TC takes a position in the asset which is positive or negative depending on the correlation. Relative to her no-tc optimal solution, the CARA-TC agent down-weights asset 2 to mitigate the impact of transaction costs.

Figure 5 illustrates the aim portfolio holdings of asset 1 (left panel) and asset 2 (right panel) as a function of κ , for constant $\sigma_{x2} = -0.5$. For asset 1, the aim portfolios are the same and constant for CMV-NoTC and CARA-NoTC agents while CMV-TC and non-myopic CARA-TC reduce their position in asset 1 as κ increases. Further, while the CARA-NoTC agent also reduces her position in asset 2, non-myopic CARA-TC agent responds more dramatically as κ is increased.¹³

In the next section, we propose an empirical application of our approach using a real world setup, to investigate in a more realistic setting the importance of hedging demands for the performance of a dynamic trading strategy.

5 Empirical Application with Retail Order Imbalance

Boehmer, Jones, Zhang, and Zhang (2021, BJZZ) 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. A trading strategy designed to take advantage of such predictability would involve large transaction costs, given the short half-life of the signals. In this section, we implement our methodology to determine optimal the optimal trading for such a strategy. We calibrate transaction costs based on a large institutional order data. We illustrate that there are economically significant utility benefits of utilizing our approach when compared to a

¹³We observe in both figures 4 and 5 that the aim portfolios for asset 1 are identical in this example for the CARA and CMV agents, with or without t-costs. While this is clear from the equations in the case without t-costs, is is not obvious in the case with t-costs and, in fact, we were unable to prove it analytically. However, it seems to hold for all parameters we tried within this admittedly very special example, where asset 1's return risk is perfectly orthogonal to the shocks to expected returns that can be perfectly hedged with asset 2. Still, it is quite remarkable.

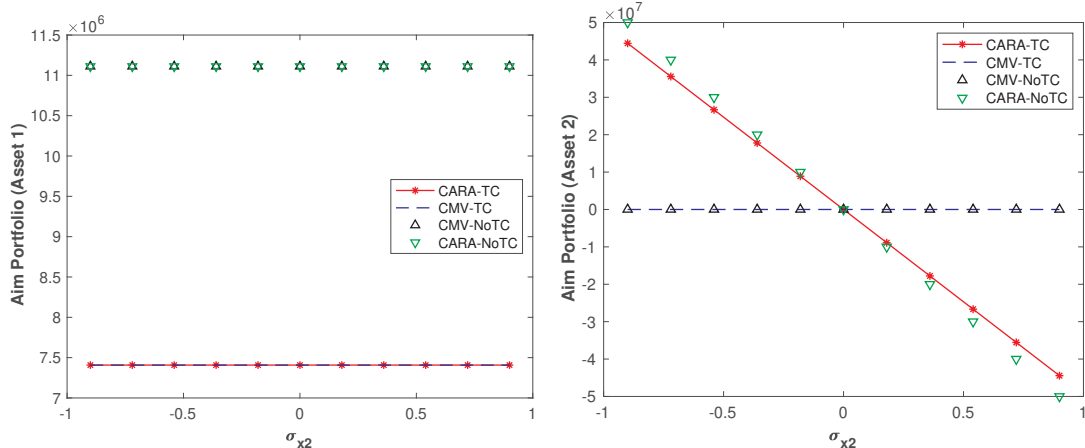


Figure 4: Aim portfolios in two assets as a function of σ_{x2} .

Notes: We plot the aim portfolios in two assets for for CMV and CARA agents in the presence and absence of transaction costs as a function of σ_{x2} , the correlation between the innovations of asset 2 and the predictor. Parameters: $\mu_0 = 0$, $\mu = 1$, $\kappa = 0.5$, $\sigma_1 = 0.3$, $\sigma_2 = 0.1$, $\sigma_x = 0.2$, $\Lambda = 0.01\Sigma$, $\gamma_x = 10^{-6}$, $\gamma = 10^{-6}$, $x = 1$, $\rho = 1$.

CMV investor who fundamentally ignores the correlation between the innovations in asset returns and predictors.

We note that, in calculating the optimal CMV and CARA portfolios we will use the full sample estimates for a number of parameters. Thus our results will overstate the extent to which we can profit from this predictability we document here. However, the utility gains our methodology provides should be valid even in a full out-of-sample exercise.

5.1 Predictability model

We begin by examining the relation between daily net retail order flow and subsequent daily returns over the period from January 2014 through December 2019.

Let $N_{i,t}^b$ ($N_{i,t}^s$) be the estimated number of retail buy (sell) trades for stock i on day t , based on the BJZZ retail trade classification algorithm. Our return predictor at the stock level is then given by

$$x_{i,t} = \frac{N_{i,t}^b - N_{i,t}^s}{N_{i,t}^b + N_{i,t}^s}.$$

Using the top 25 largest market capitalization US common stocks as of January 1, 2014, we estimate the following stock-level regressions:

$$r_{i,t+1} = \beta_{0,i} + \beta_i x_{i,t} + \epsilon_{i,t+1} \quad (77)$$

$$\Delta x_{i,t+1} = \phi_{0,i} - \phi_i x_{i,t} + \varepsilon_{i,t+1} \quad (78)$$

Table 1 summarizes the regression results. We find that in 21 out of 25 regressions, $\hat{\beta}_i$ is positive. β_i is statistically significant at 1% (5%) level for three (six) stocks. More importantly, the mean

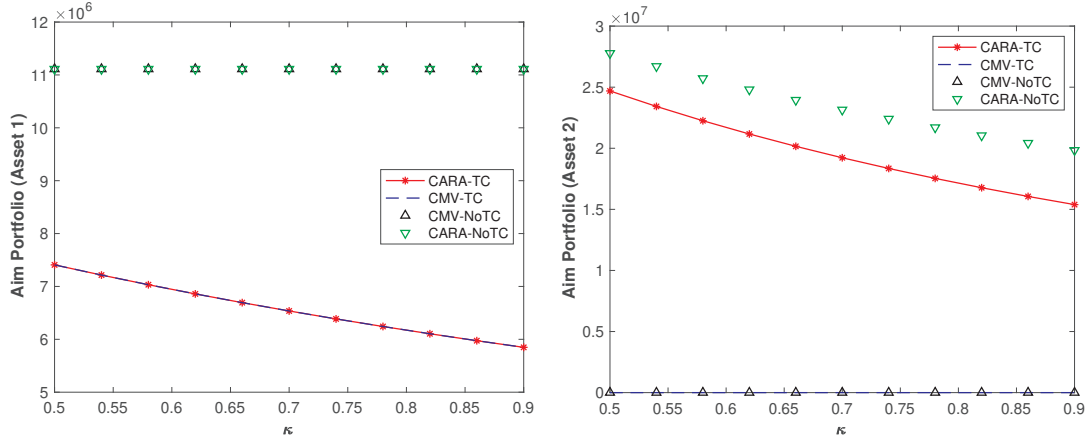


Figure 5: Aim portfolios in two assets as a function of κ .

Notes: We plot the aim portfolios in two assets for for CMV and CARA agents in the presence and absence of transaction costs as a function of κ , the mean reversion speed of the predictor. We use the following parameters: $\mu_0 = 0$, $\mu = 1$, $\sigma_1 = 0.3$, $\sigma_2 = 0.1$, $\sigma_x = 0.2$, $\sigma_{x2} = -0.5$, $\Lambda = 0.01\Sigma$, $\gamma_x = 10^{-6}$, $\gamma = 10^{-6}$, $x = 1$, $\rho = 1$.

estimate for β is highly statistically significant: $\widehat{\beta}_i = 0.0031$, with a corresponding standard error of 0.0005. For the majority of the stocks, $\beta_{0,i}$ is estimated to be statistically insignificant than zero.

For each of the 25 regressions, $\widehat{\phi}_i$ is statistically significantly different from both zero and one at 1% level. Specifically, the mean estimate is 0.45 with a corresponding standard error of 0.018. This estimate corresponds to a half life of 1.2 trading days.

Table 1 also reports the estimated correlation between return innovations ($\epsilon_{i,t+1}$) and innovations in the predictive variable ($\epsilon_{i,t+1}$) each stock (conditional on information up to t). The estimated correlation is negative for every stock in the sample, and ranges from -0.13 to -0.42. These estimates is consistent with a contrarian trading strategy at the aggregate retail level: on average, retail traders tend to sell (buy) a stock with positive (negative) daily excess returns. Overall, despite relatively high mean-reversion rates, this analysis illustrates that this predictability model is very appropriate for our framework with hedging demands.

5.2 Average utility and calibration of main parameters

We now run various experiments with different number of stocks (from the list in Table 1) and different time horizons to quantify the utility benefits of following the optimal trading policy derived for a non-myopic CARA agent in the presence of trading costs when compared to the policy derived for a CMV agent. We use the finite horizon model developed in Section 2, and trade at the end of each day based on the (daily) retail-order imbalance measure.

For an experiment with a finite horizon of T , we will compute average utility of each policy from the maximum available non-overlapping samples between 2014 and 2019. For example, if the investment horizon, T , is three months, then there will be 24 samples.

Let $w_{t,k}^j$ be the wealth of the agent j (CARA or CMV) at time t corresponding to sample k in

the experiment. Then, the utility of each agent in this experiment is given by $U_k^j = -e^{-\gamma w_{T,k}^j}$ and the average utility across all sample paths is given by

$$U_{\text{avg}}^j = \frac{1}{K_T} \sum_{k=1}^{K_T} U_k^j,$$

where K_T denotes the number of samples corresponding to an horizon of T .¹⁴

Finally, the total wealth net of transaction costs in sample k for an agent with type j is given by

$$w_{T,k}^j = \sum_{t=1}^T \left[\left(x_{t,k}^j \right)^\top \Delta S_{t,k} - \frac{1}{2} \left(\Delta x_{t,k}^j \right)^\top \Lambda \left(\Delta x_{t,k}^j \right) \right]$$

where $\Delta S_{t,k}$ is the vector of realized price changes at time t and $x_{t,k}^j$ is the vector of position holdings for an agent with type j in sample path k .

Since both CMV and CARA agents use price-changes for asset returns, we now discuss the calibration of the main inputs to the agent policies at the beginning of each sample path k . The parameters are based on the statistics that we estimate from the full period.

Let V be the estimated full-sample variance-covariance matrix of the daily returns given by $V = \text{Var}(\epsilon)$. Since both CARA and CMV need the variance of price changes, we use

$$\Sigma(k) = \text{diag}(S_{0,k}) V \text{diag}(S_{0,k})$$

where $S_{0,k}$ is the price vector of the stocks at the beginning of sample k . We set $\mu_0(k)$ to be the zero vector as the intercept term in regression specified in equation (77) is estimated to be insignificant from zero for the majority of the stocks. Finally, to convert to a price-change-based expected return, $\mu(k)$ will equal the β scaled by the initial stock prices:

$$\mu(k) = \text{diag}(S_{0,k}) \text{diag}(\beta)$$

The parameters for the predictor dynamics do not require scaling. We set $\kappa = \text{diag}(\hat{\phi})$ and $\Sigma_x = \text{Var}(\hat{\epsilon})$. These are held constant for every sample path k . Finally, $\Sigma_{sx}(k) = \text{diag}(S_{0,k}) \text{Cov}(\epsilon, \epsilon)$.¹⁵

5.3 Calibration of the Price Impact Matrix

To calibrate the transaction cost multipliers of our model realistically, we use proprietary execution data from the historical order databases obtained from 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

¹⁴There are 1510 trading days in this six year horizon so $K_T = \lfloor \frac{1510}{T} \rfloor$.

¹⁵ $\sigma_x(k)$, $\sigma_{xs}(k)$ and $\Omega(k)$ can be derived from these definitions.

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%. This data set has been utilized in Bogousslavsky, Collin-Dufresne, and Sağlam (2021) to study the impact of trading glitch at Knight Capital on institutional trading costs. We filter this data set by focusing on large-order trades on the same set of 25 stocks listed in Table 1. The data set has a large order execution on all of these stocks except one (BRKB). There are 9,405 large-order executions on this subset.

A standard measure of institutional trading costs is the implementation shortfall (IS), which is defined 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 i th parent-order is given by

$$IS_i = D_i \frac{P_i^{\text{avg}} - P_{i,0}}{P_{i,0}}, \quad (79)$$

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) at the point in time when the parent order execution begins.

We estimate the price impact coefficient θ by running the regression

$$IS_i = \alpha_i + \theta \frac{Q_i}{Vlm_i} + \varepsilon_i \quad (80)$$

over all orders for our 25 firms. Here, Q_i is the number of shares bought or sold, and Vlm_i is the daily volume of the stock realized during the execution day. Thus, $\frac{Q_i}{Vlm_i}$ measures the size of the order as a fraction of the daily market volume.

Table 2 reports that θ is estimated to be 0.0465 (465 basis points). The reported standard errors are double-clustered at the stock-day level. We find that θ is statistically significant at 1% level. The economic magnitude is also large. That is an order that traded 1% of daily volume resulted in a transaction cost of 4.65 bps.

According to our quadratic transaction cost model, trading $u_{i,t}$ shares of stock i on day t would move its (average) price by $\frac{\lambda_i u_{i,t}}{2}$ where λ_i is the stock's price impact coefficient. On average, our empirical model based on implementation shortfall would predict this price impact to be $\theta u_{i,t} \bar{z}_i$ where \bar{z}_i is the average ratio of the stock's price to its volume, i.e., $\frac{1}{T_{\max}} \sum_{t=1}^{T_{\max}} \frac{S_{i,t}}{V_{i,t}}$ where $V_{i,t}$ is the daily volume of stock i on day t and T_{\max} is equal to the number of days in our six-year horizon (1510 trading days). According to this relation, $\lambda_i = 2\theta \bar{z}_i$. The final column in Table 1 reports the price impact coefficient for each stock in our sample. The highest price impact of trading one share

is observed on AMZN, partially attributable to its high share price during our sample period.

For simplicity, we assume that there is no cross-price impact. Thus, we calibrate Λ to be $\text{diag}(\lambda)$ for every sample k .¹⁶

5.4 Insights from the one-asset and one-predictor experiments

To gain a better understanding of the differences between the trading policies of CARA and CMV agents, we first examine the results from one-asset and one-predictor example for each stock in our sample. We use the first sample ($k = 1$) corresponding to three-month horizon ($T = 63$) and set $\gamma = 10^{-9}$ which we can think of as corresponding to a relative risk aversion of 1 for an agent with 1 billion dollars under management. We compute the aim portfolio at $t = 0$ assuming that the predictor is equal to 1, i.e., $x_{i,t} = x_{i,0} = 1$.

Figure 6 illustrates the percentage deviation of the CARA aim portfolio from the CMV aim portfolio as a function of the estimated model parameters.¹⁷ Since the correlation between return and factor innovations is negative for all stocks in our sample, we find that the aim portfolio of the CARA agent is more aggressive when compared to that of the CMV agent. The percentage difference is higher than 5% for T, XOM and VZ suggesting that there are economically significant deviations from the CMV portfolio.

The top-left plot shows that the percentage difference is overall increasing in the Sharpe ratio of the predictor as proxied by $\frac{\mu}{\sigma_s}$. When you consider the relation between t-statistics and Sharpe ratio in the return predictability model, this also partially explains why the percentage difference in aim portfolios is the highest for T, XOM and VZ as these stocks' β estimates have the highest t-statistics. This result is very striking suggesting that the Sharpe ratio of the predictor drives the main difference in the resulting aim portfolios.

The top-right plot shows that the percentage difference is overall decreasing in the mean reversion speeds of the predictor, κ . When the signal is more persistent, the hedging demand term becomes more significant. Similarly, the bottom-left plot shows the CARA agent has higher hedging demand when the signal innovations has more negative correlation with the return innovations. Finally, for the existing sample, λ is not highly correlated with the difference in the aim portfolios of the CMV and CARA agents, as shown in the bottom-right panel.

Figure 7 illustrates the percentage difference between the trading speeds of the CARA and CMV agents as a function of the estimated model parameters. Since the correlation between return and factor innovations is negative for all stocks in our sample, we find on average that the CARA agent trades slowly to the target aim portfolio to compensate the additional position size. We find that the relative trading speed of the CARA agent decreases as the Sharpe ratio or the persistence of the signal is higher, and the correlation between the predictor and return innovations is more negative.

¹⁶Our results are very similar if we define $\bar{z}_i = \frac{1}{T_{\max}} \sum_{t=1}^{T_{\max}} \frac{1}{V_{i,t}}$ and $\Lambda(k) = \text{diag}(S_{0,k}) \text{diag}(\lambda)$ for every sample path k .

¹⁷Formally, the percentage deviation is given by $100 \left(\frac{\text{aim}(\text{CARA-TC})}{\text{aim}(\text{CMV-TC})} - 1 \right)$. Note that the aim portfolio for the CARA and CMV agent is negative for AAPL, PG, PFE and BAC as the estimated β is negative for these stocks as shown in Table 1.

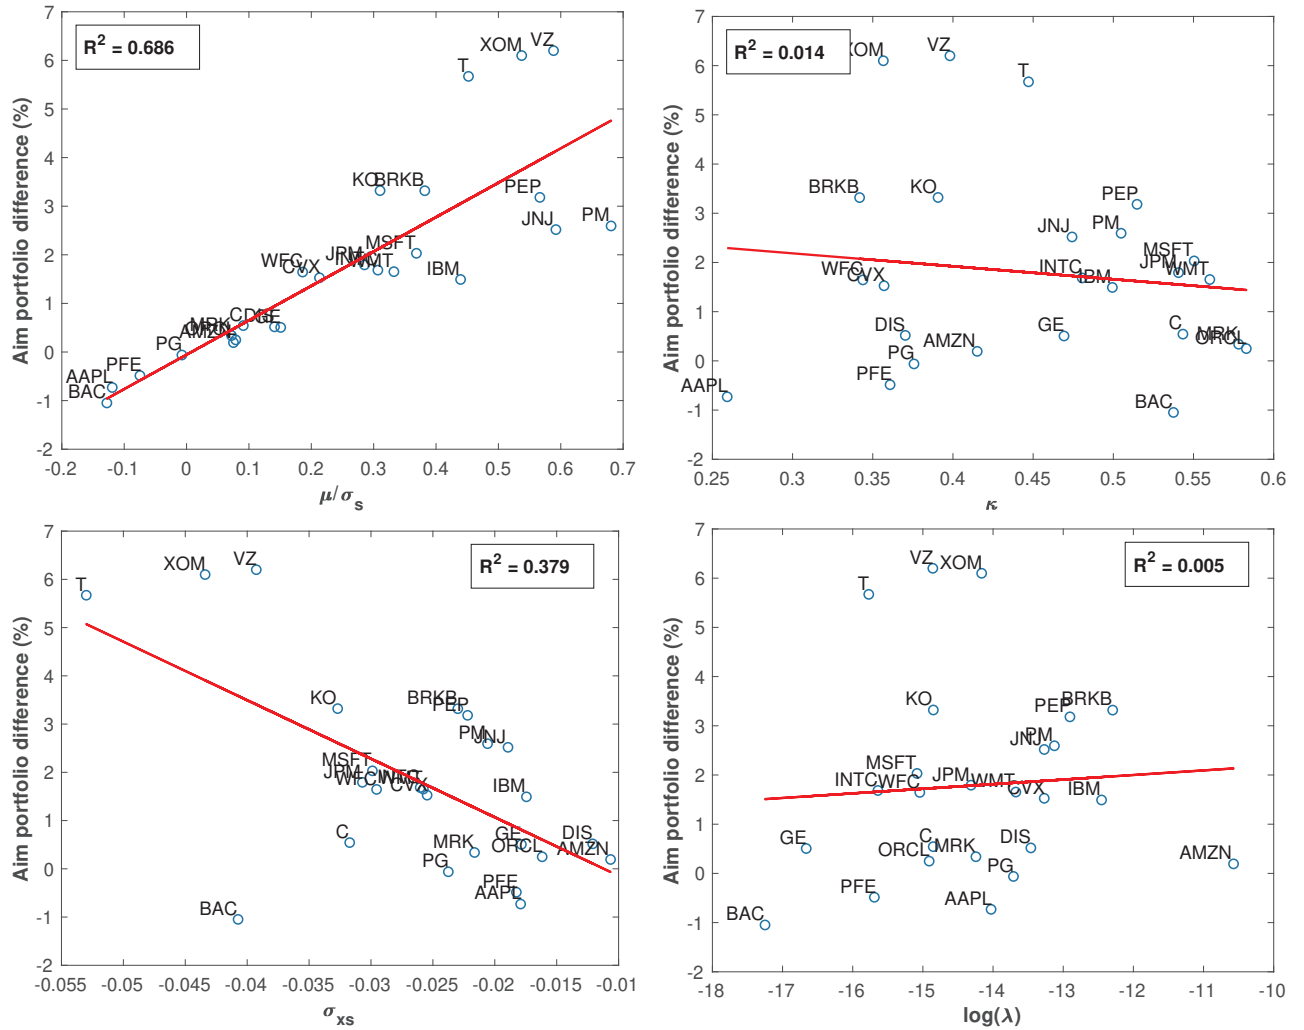


Figure 6: The percentage differences between the aim portfolios of the CARA and CMV agents in the one-asset and one-predictor experiment as a function of the estimated model parameters.

Notes: Formally, the percentage deviation is given by $100 \left(\frac{\text{aim}(\text{CARA-TC})}{\text{aim}(\text{CMV-TC})} - 1 \right)$. Note that the aim portfolio for the CARA and CMV agent is negative for AAPL, PG, PFE and BAC as the estimated β is negative for these stocks.

We again find that the main driver in the differences in trading speeds is the Sharpe ratio of the asset.

5.5 Numerical experiments with multiple assets and predictors

In this section, we report the results of numerical experiments with various number of assets ($N = 2$, $N = 5$ and $N = 25$), horizon (three-month, $T = 63$, and six-month, $T = 126$) and risk aversion level. We use PFE and T (JPM, PM and T) [JNJ, JPM, KO, PG, and VZ] in the two-asset (three-asset) [five-asset] experiments and use all the stocks in Table 1 in the case of $N = 25$. We consider low and high risk aversion cases corresponding to $\gamma = 5 \times 10^{-9}$ and $\gamma = 10 \times 10^{-9}$, respectively. In total, we assess the performance of CARA and CMV trading policies in 16 different experiments. Table 3 summarizes these by reporting the average utility, certainty equivalent wealth and the Sharpe ratios across all experiments.¹⁸

The first row of each experiment reports the average utilities achieved by the CARA and CMV agents and the difference in the utilities. We test whether the difference is statistically different than 0 by performing a two-tailed t -test. We find that in 14 out of 16 experiments, the difference is statistically significant at 5% level indicating that there are statistically significant utility benefits in following the CARA trading policy. In the second row of each experiment, we report the certainty equivalent wealth ($CE(w_T)$) achieved by each policy where

$$CE(w_T) = -\frac{1}{\gamma} \log \left(\frac{1}{K_T} \sum_{k=1}^{K_T} -e^{-\gamma w_{T,k}} \right).$$

In the third column, we report the percentage improved achieved by the CARA trading policy. The difference in utility is highly significant in the four experiments with $N = 3$. Note that this is expected as this set includes PM and T, which have high expected returns with low risk. We find that the improvement in the certainty equivalent can reach as high as 7% indicating economically significant benefits. In the third row of each experiment, we report the Sharpe ratios and the percentage improvement achieved by the CARA agent. Since the Sharpe ratio of the terminal wealth can also approximate our objective function (despite ignoring higher moments), we find consistently that the CARA agent achieves higher Sharpe ratio in all experiments with improvements up to 3%.

For robustness, we have also run 2,300 three-asset experiments (25 choose 3 combinations) for three-month investment horizon and high-risk aversion. We find that the mean improvement in $CE(w_T)$ is 1.9% with a standard error of 0.13%. Further, we examine the top 100 portfolios having the highest improvement over the CARA agent with regard to CE , and for this group the average improvement goes up to 13%. In these portfolios, the most frequent stock is T with 84 occurrences out of 300. This is anticipated partly due to its lowest σ_{xs} value across all stocks along with its relatively high reward-to-risk ratio.

¹⁸Note that since the terminal wealth is not normally distributed, the Sharpe ratio does not align perfectly with our objective function but we report it as it is a first-order measure of risk-adjusted performance in practice.

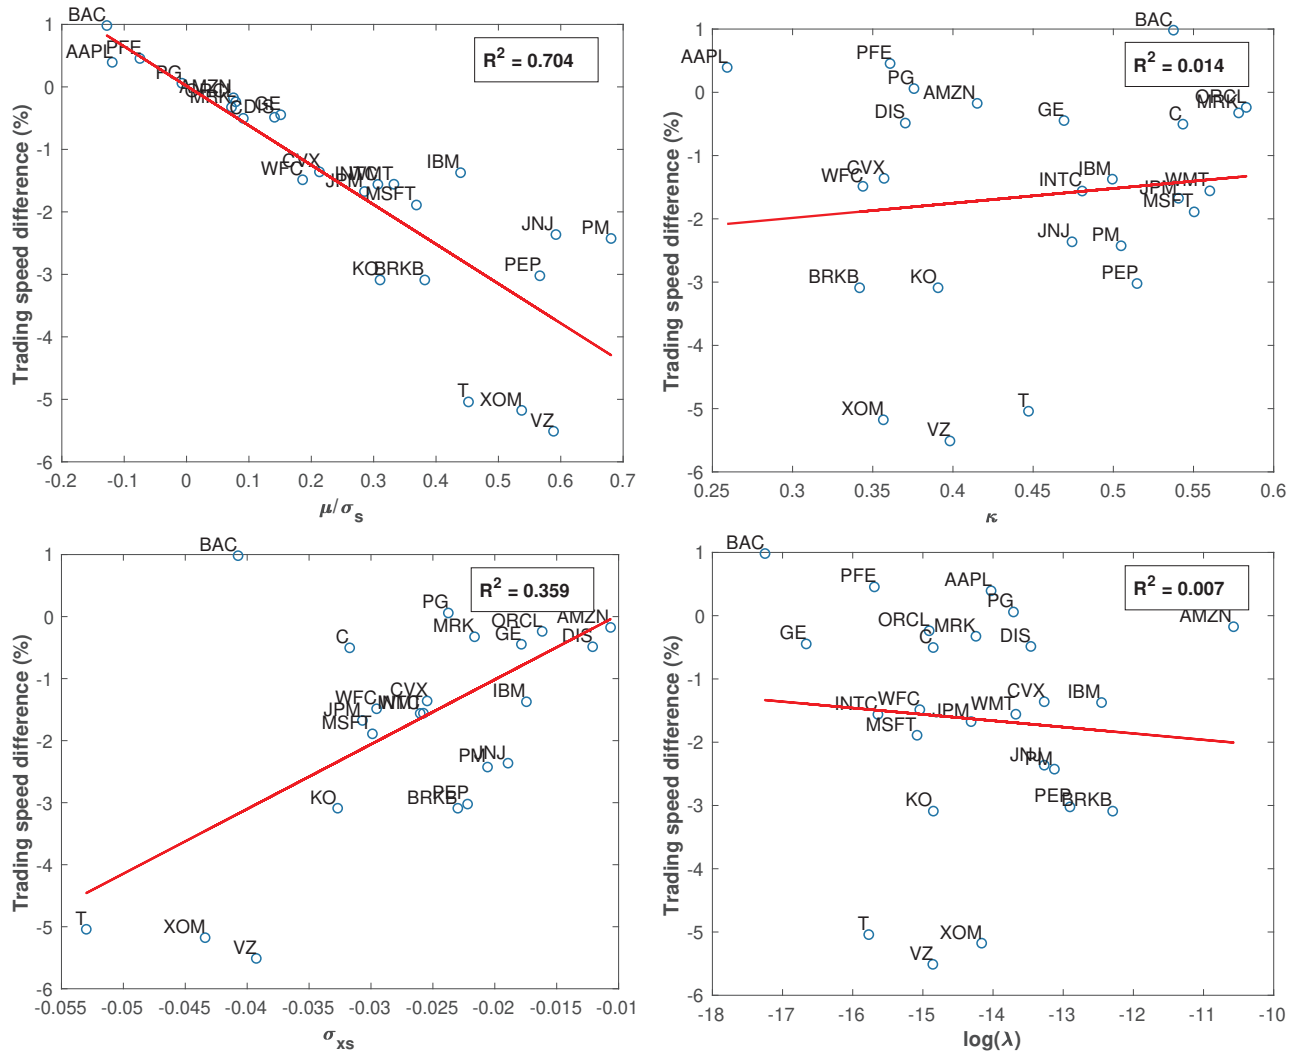


Figure 7: The differences between the trading speeds of the CARA and CMV agents in the one-asset and one-predictor experiment as a function of the estimated model parameters.

Notes: Formally, the percentage deviation in trading speeds is given by $100 \left(\frac{\tau(\text{CARA-TC})}{\tau(\text{CMV-TC})} - 1 \right)$.

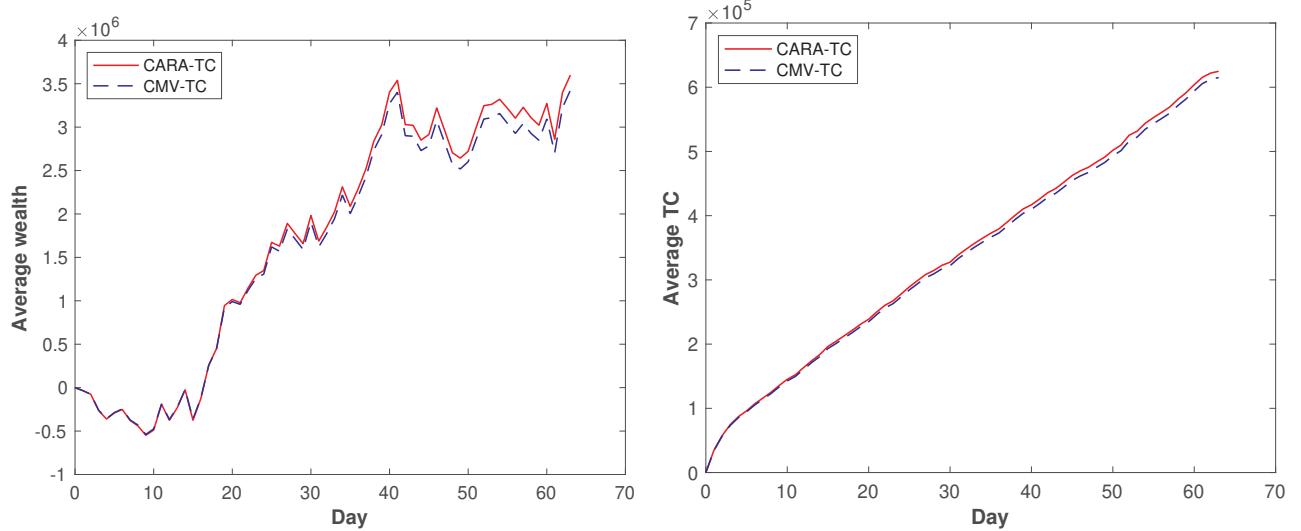


Figure 8: Average wealth and cumulative transaction cost on each day from the five-asset experiment with three-month horizon and high risk aversion. The five asset experiment includes JNJ, JPM, KO, PG, and VZ in the portfolio.

To further examine the trading policies in the multi-asset experiments, Figure 8 plots the average cumulative wealth and transaction costs on each day in the five-asset experiment with three-month horizon (i.e., 24 samples) and high risk aversion. Beginning approximately on day 25, the average wealth of the CARA agent becomes significantly higher than that of the CMV agent. Overall, total trading costs roughly account for 20% of the average wealth achieved for both agents, but compared to the CMV agent, we find that the CARA agent pays slightly higher total trading costs. Despite the potential lower trading speeds, the CARA agent may hold larger aim portfolio and end up paying higher trading costs. We examine these issues in detail by focusing on the differences in trading on a given single sample. Figure 9 plots the predictor, position and trading speeds (the diagonal entries in the matrix) on PG and VZ on a single path. Here, we focus on PG and VZ as compared to PG, VZ has higher Sharpe ratio, lower mean reversion speed and larger negative correlation between its return and predictor innovations. Given these differences between two assets, we find that the CMV and the CARA agent have roughly the same position and trading speed in PG but they differ a lot in VZ. Specifically, the CARA agent has significantly larger position in VZ and lower trading speed when compared to the CMV agent. Compared to the one-asset example, these trading policies are also driven by other assets' predictors and correlation structure but overall we observe that our earlier insights extend into these multi-asset and multi-predictor experiments as well.

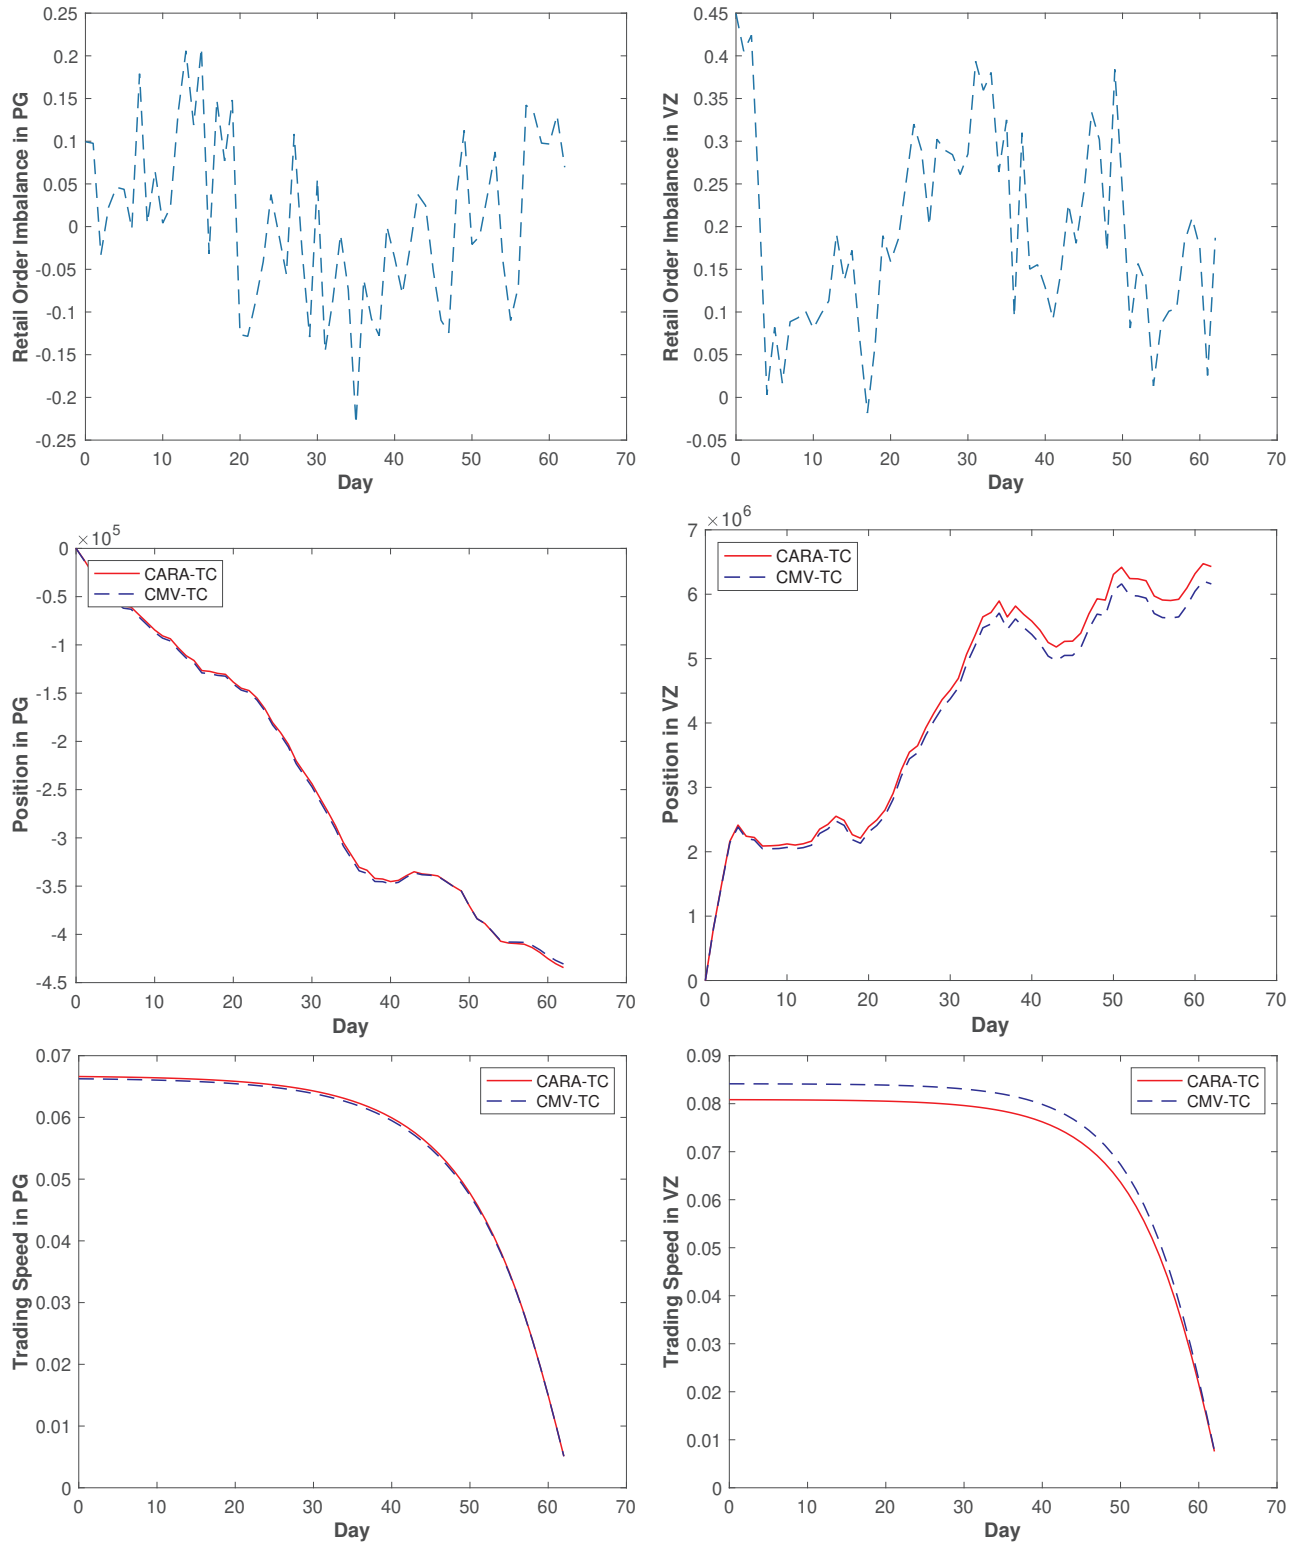


Figure 9: Predictor, position and trading speeds on PG and VZ in the five-asset experiment with three-month horizon and high risk aversion.

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 utilize the retail order imbalance as a predictor at the stock-level, we find that hedging demands significantly affect strategy performance.

Rank	Ticker	$\hat{\beta}_i$	$t(\beta_i)$	$\hat{\phi}_i$	$t(\phi_i)$	$\text{cor}(\hat{\epsilon}_i, \hat{\epsilon}_i)$	$\hat{\lambda}_i(\times 10^7)$
1	AAPL	-0.0018	-0.59	0.26	11.27	-0.21	8.09
2	XOM	0.0063	2.64	0.36	12.44	-0.41	7.07
3	MSFT	0.0053	1.37	0.55	19.12	-0.26	2.82
4	GE	0.0027	0.69	0.47	12.90	-0.14	0.58
5	JNJ	0.0060	2.09	0.47	13.91	-0.22	17.25
6	WMT	0.0039	1.51	0.56	15.13	-0.23	11.50
7	CVX	0.0029	0.80	0.36	10.91	-0.31	17.25
8	WFC	0.0024	0.87	0.34	14.42	-0.27	2.93
9	JPM	0.0037	0.94	0.54	18.00	-0.32	6.08
10	PG	-0.0001	-0.04	0.38	13.27	-0.24	11.12
11	IBM	0.0055	1.32	0.50	14.68	-0.25	39.04
12	PFE	-0.0008	-0.42	0.36	10.92	-0.16	1.53
13	T	0.0050	2.61	0.45	17.60	-0.42	1.41
14	AMZN	0.0014	0.18	0.42	12.55	-0.18	256.15
15	KO	0.0028	1.42	0.39	13.36	-0.32	3.55
16	BAC	-0.0020	-0.59	0.54	17.75	-0.30	0.32
17	ORCL	0.0010	0.27	0.58	15.45	-0.13	3.35
18	C	0.0014	0.44	0.54	17.57	-0.27	3.54
19	MRK	0.0009	0.31	0.58	17.77	-0.21	6.51
20	VZ	0.0062	3.16	0.40	13.89	-0.38	3.52
21	BRKB	0.0039	1.57	0.34	11.13	-0.29	45.75
22	PM	0.0086	2.50	0.50	15.00	-0.21	19.94
23	DIS	0.0017	0.53	0.37	10.75	-0.14	14.25
24	INTC	0.0048	1.27	0.48	17.30	-0.24	1.61
25	PEP	0.0051	2.06	0.51	18.39	-0.24	24.88

Table 1: For the 25 largest US common stocks by equity market capitalization as of January 1, 2014, we estimate equations (77) and (78) at the individual firm level using daily net retail order flow and subsequent daily returns. We also estimate the price impact coefficient using an institutional large order data set. Standard errors are adjusted for serial correlation and heteroscedasticity using Newey-West estimator with a maximum lag of ten days. The calculation of the individual-firm price impact parameter $\hat{\lambda}_i$ is described in detail in Section 5.3.

	<i>Dependent variable:</i>
	IS (bps)
θ (bps)	465.14*** (95.56)
Constant (α)	0.58 (1.14)
Observations	9,405
Adjusted R ²	0.001

Table 2: Estimation of the price impact coefficient, θ , from institutional order data set
Notes: We estimate the regression model specified in equation (80) using 9,405 large-order executions on the same 25-stock universe. The reported standard errors are clustered at the day and stock level.

		<i>Low risk aversion</i> $\gamma = 5 \times 10^{-9}$			<i>High risk aversion</i> $\gamma = 10 \times 10^{-9}$		
(N, T)	Statistic	CARA-TC (1)	CMV-TC (2)	Diff (3)	CARA-TC (1)	CMV-TC (2)	Diff (3)
$N = 2$	U_{avg}	-0.9781	-0.9794	0.0013***	-0.9746	-0.9763	0.0017***
$T = 63$	$CE(w_T)$ (\$M)	4.43	4.17	6.4%	2.57	2.40	7.1%
	Sharpe ratio	1.13	1.11	1.5%	1.01	0.98	2.7%
$N = 2$	U_{avg}	-0.9714	-0.9732	0.0018*	-0.9666	-0.9688	0.0022*
$T = 126$	$CE(w_T)$ (\$M)	5.79	5.43	6.7%	3.40	3.17	7.1%
	Sharpe ratio	0.82	0.80	1.0%	0.77	0.76	1.8%
$N = 3$	U_{avg}	-0.9643	-0.9658	0.0015***	-0.9535	-0.9556	0.0020***
$T = 63$	$CE(w_T)$ (\$M)	7.27	6.97	4.4%	4.76	4.54	4.7%
	Sharpe ratio	1.83	1.80	1.4%	1.63	0.59	2.5%
$N = 3$	U_{avg}	-0.9408	-0.9431	0.0023***	-0.9225	-0.9256	0.0031***
$T = 126$	$CE(w_T)$ (\$M)	12.20	11.70	4.2%	8.07	7.73	4.4%
	Sharpe ratio	1.19	1.16	2.5%	1.10	1.07	2.8%
$N = 5$	U_{avg}	-0.9772	-0.9782	0.0010**	-0.9672	-0.9687	0.0015**
$T = 63$	$CE(w_T)$ (\$M)	4.61	4.40	4.7%	3.34	3.18	5.0%
	Sharpe ratio	0.97	0.95	1.6%	0.97	0.95	2.1%
$N = 5$	U_{avg}	-0.9654	-0.9670	0.0017**	-0.9494	-0.9519	0.0025**
$T = 126$	$CE(w_T)$ (\$M)	7.05	6.71	5.1%	5.19	4.93	5.4%
	Sharpe ratio	0.67	0.65	2.7%	0.67	0.65	2.9%
$N = 25$	U_{avg}	-0.8729	-0.8753	0.0024**	-0.8370	-0.8406	0.0036***
$T = 63$	$CE(w_T)$ (\$M)	27.19	26.65	2.0%	17.79	17.36	2.5%
	Sharpe ratio	1.00	0.98	1.9%	1.06	1.04	2.1%
$N = 25$	U_{avg}	-0.7627	-0.7667	0.0040**	-0.7003	-0.7059	0.0056**
$T = 126$	$CE(w_T)$ (\$M)	54.19	53.14	2.0%	35.63	34.83	2.3%
	Sharpe ratio	1.00	0.98	1.7%	0.67	0.65	1.9%

Table 3: Policy comparison across 16 numerical experiments with various number of assets ($N = 2$, $N = 3$, $N = 5$ and $N = 25$), horizon (three-month, $T = 63$, and six-month, $T = 126$) and risk aversion level.

References

- Almgren, Robert, and Neil Chriss, 2001, Optimal execution of portfolio transactions, *Journal of Risk* 3, 5–40.
- 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.
- Bogousslavsky, Vincent, Pierre Collin-Dufresne, and Mehmet Sağlam, 2021, Slow-moving capital and execution costs: Evidence from a major trading glitch, *Journal of Financial Economics* 139, 922–949.
- 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.

- Delbaen, F, Y Hu, and A Richou, 2011, On the uniqueness of solutions to quadratic bsdes with convex generators and unbounded terminal conditions, *Annals of the Institute Henri Poincare in Probability and Statistics* 47, 559–574.
- 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.
- Kobylanski, Magdalena, 2000, Backward stochastic differential equations and partial differential equations with quadratic growth, *Annals of Probability* 28, 558–602.
- 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.
- Muhle-Karbe, Johannes, James A Sefton, and Xiaofei Shi, 2023, Dynamic portfolio choice with intertemporal hedging and transaction costs, *Available at SSRN 4522752*.
- 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 $M_t = E_t[U(W_T)]$ 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

$$dH_t = \mu_H dt + \sigma_H^\top dZ_t \tag{81}$$

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:

$$\mu_H = -\frac{1}{2} \frac{U''(H)}{U'(H)} \|\sigma_H\|^2 \tag{82}$$

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:

$$H_t = E_t[W_T - \int_t^T \mu_H(H_t, \sigma_H) dt] \tag{83}$$

where the driver of the BSDE is given in equation (82) above.

To summarize, we have shown that, for an agent with an arbitrary wealth process W_t driven by a vector of N Brownian motions, who has expected utility of terminal wealth $E[U(W_T)]$, we can define his certainty equivalent H_t in two different ways. First, the traditional definition $H_t = U^{-1}(E_t[U(W_T)])$. Second, as the solution of the BSDE given in (82-83) above. Both are equivalent.¹⁹ 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.

¹⁹We note that this also provides a direct constructive proof for the uniqueness of the solution to that particular quadratic BSDE, a subject studied in more general settings in Kobylanski (2000) or Delbaen, Hu, and Richou (2011).

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:

$$H_t = \mathbb{E}_t[W_T - \int_t^T \mu_H(H_t, \sigma_{H,s}, \sigma_{H,x}) dt] \quad (84)$$

$$\mu_H = -\frac{1}{2} \frac{U_1''(H)}{U_1'(H)} \|\sigma_{H,s}\|^2 - \frac{1}{2} \frac{U_2''(H)}{U_2'(H)} \|\sigma_{H,x}\|^2 \quad (85)$$

where two different (twice-differential, strictly increasing and concave) utility functions U_i $i = 1, 2$ apply to the different sources of diffusion risk. 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:

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

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:

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

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

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

which is, indeed, the objective function we consider in equations (93) and (101) 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 $\gamma_1 = \gamma$):

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

Now, let's instead consider a CMV-agent that, for general x_t , maximizes the objective function:

$$J(x_t, n_t, t) = \max_{\theta_t} \mathbb{E}_t \left[\int_t^T \left\{ n_u^\top (\mu_0 + \mu x_u) du - \frac{1}{2} \theta_u^\top \Lambda \theta_u - \frac{1}{2} \gamma n_u^\top \Sigma n_u \right\} du \right] \quad s.t. \quad dn_t = \theta_t dt \quad (86)$$

The HJB for that problem is as follows:

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

This equation is identical to the HJB equation (105) of the recursive-utility-source-dependent agent given in appendix F below when substituting the same Σ_x , but setting $\gamma_x = 0$ and $\sigma_{xs} = 0$ (and thus $\Omega = 0$). Further, the solution of this equation discussed in detail in G below shows that portfolio strategy of the CMV agent is independent of σ_x (as well as σ_{sx}). Whence the corollary.

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:

$$dH_t = \mu_H dt + \sigma_{H,s} dZ^s + \sigma_{H,x} dZ^x + \eta_H (dN_t - \rho dt). \quad (87)$$

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

$$\mathcal{U}(H_t, 0, 0, 0) = \mathbb{E}_t[\mathcal{U}(H_t + \mu_H dt, \sigma_{H,s} dZ^s, \sigma_{H,x} dZ^x, \eta_H (dN_t - \rho dt))] \quad (88)$$

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 (88), 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:

$$\mu_H = -\frac{1}{2}\frac{U_1''(H)}{U_1'(H)}\|\sigma_{H,s}\|^2 - \frac{1}{2}\frac{U_2''(H)}{U_2'(H)}\|\sigma_{H,x}\|^2 + \rho \left(\eta_H - \frac{U_3(H_t + \eta_H) - U_3(H_t)}{U_3'(H_t)} \right) \quad (89)$$

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

$$H_t = \mathbb{E}_t \left[W_T - \int_t^T \left\{ \frac{1}{2}\gamma_1\|\sigma_{H,s}\|^2 + \frac{1}{2}\gamma_2\|\sigma_{H,x}\|^2 + \rho \left(\eta_H - \frac{1 - e^{-\gamma_3 \eta_H}}{\gamma_3} \right) \right\} ds \right] \quad (90)$$

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 (84). 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 $t \leq \mathcal{T}$:

$$\begin{aligned} H_t = & \mathbb{E}_t \left[W_{\mathcal{T}} - \int_t^{\mathcal{T}} \left\{ \frac{1}{2}\gamma_1\|\sigma_{H,s}\|^2 + \frac{1}{2}\gamma_2\|\sigma_{H,x}\|^2 + \rho \left(W_s - H_{s-} - \frac{1 - e^{-\gamma_3(W_s - H_{s-})}}{\gamma_3} \right) \right\} ds \right] \\ = & W_t \\ & + \mathbf{1}_{\{\mathcal{T} > t\}} \mathbb{E}_t \left[\int_t^{\infty} e^{-\rho(s-t)} \left\{ dW_s - \left[\frac{1}{2}\gamma_1\|\sigma_{H,s}\|^2 + \frac{1}{2}\gamma_2\|\sigma_{H,x}\|^2 + \rho \left(W_s - H_{s-} - \frac{1 - e^{-\gamma_3(W_s - H_{s-})}}{\gamma_3} \right) \right] ds \right\} \right] \end{aligned} \quad (91)$$

The equality between the first and second line requires an additional transversality condition.²⁰

²⁰Note that $\mathbb{E}_t \left[\int_t^{\mathcal{T}} dX_u \right] = \mathbb{E}_t \left[\int_t^{\infty} \rho e^{-\rho(s-t)} ds \int_t^s dX_u \right] = \mathbb{E}_t \left[\int_t^{\infty} e^{-\rho(s-t)} dX_s - [e^{-\rho(s-t)}(X_s - X_t)]_t^{\infty} \right]$. There-

We prove equation 33 of theorem 6.

Theorem 6

When $\gamma_1 = \gamma_2 = \gamma_3 = \gamma$ then the solution to (91) is the certainty equivalent of a CARA investor with expected utility of terminal wealth generated at the random horizon \mathcal{T} . That is $H_t = \frac{1}{\gamma} \log(\mathbb{E}_t[e^{-\gamma W_{\mathcal{T}}}]$.

Proof. Note that the solution to (91) 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 $\mathcal{T} > t$:

$$dH_t = \mu_H dt + \sigma_{H,s} dZ^s + \sigma_{H,x} dZ^x + (W_t - H_{t-})(d\mathbf{1}_{\{\mathcal{T} \leq t\}} - \rho dt) \tag{92}$$

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:

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

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} . It follows that at $t < \mathcal{T}$ and using the optional stopping theorem:

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

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 $\sigma_{H,s} = n^\top \sigma_s + J_x^\top \sigma_{x_s}$ and $\sigma_{H,x} = J_x^\top \sigma_x$. It follows from equation (7) that the function $J(x, t)$ must satisfy:

fore the transversality condition is $\lim_{T \rightarrow \infty} \mathbb{E}[e^{-\rho T} X_T] = 0$.

$$J(x_t, t) = \max_n \mathbb{E}_t \left[\int_t^T \left\{ dW_u - \frac{1}{2} \gamma n_u^\top \Sigma n_u du - \frac{1}{2} J_x^\top \Omega J_x du - \gamma n_u^\top \Sigma_{sx} J_x du \right\} \right] \quad (93)$$

where we define:

$$\Omega = \gamma \sigma_{xs} \sigma_{xs}^\top + \gamma_x \sigma_x \sigma_x^\top \quad (94)$$

$$\Sigma_{sx} = \sigma_s \sigma_{xs}^\top \quad (95)$$

The corresponding Bellman-equation is:

$$0 = \max_n \mathbb{E}_t [dW_t - \frac{1}{2} \gamma n_t^\top \Sigma n_t dt - \frac{1}{2} J_x^\top \Omega J_x dt - \gamma n_t^\top \Sigma_{sx} J_x dt + dJ(t, x_t)] \quad (96)$$

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

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

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 , is²¹

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

Plugging back into the HJB equation we get:

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

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 $c_2(t)$ is a symmetric matrix, $c_1(t)$ is a K -dimensional vector and $c_0(t)$ a scalar (all deterministic functions of time). Thus:

$$J_t = \dot{c}_0 + \dot{c}_1^\top x + \frac{1}{2} x^\top \dot{c}_2 x$$

$$J_x = c_1 + c_2 x$$

$$J_{xx} = c_2$$

²¹The second order condition $\gamma \Sigma > 0$ is always satisfied.

Plugging into the HJB we obtain:

$$-\dot{c}_0 - \dot{c}_1^\top x - \frac{1}{2} x^\top \dot{c}_2 x = \frac{1}{2} (\mu_0 + \mu x - \gamma \Sigma_{sx} (c_1 + c_2 x))^\top (\gamma \Sigma)^{-1} (\mu_0 + \mu x - \gamma \Sigma_{sx} (c_1 + c_2 x)) \\ - \frac{1}{2} (c_1 + c_2 x)^\top \Omega (c_1 + c_2 x) - (c_1 + c_2 x)^\top \kappa x + \frac{1}{2} \text{Tr}(c_2 \Sigma_x)$$

This equation is satisfied provided c_0, c_1, c_2 solve the following system:²²

$$-\dot{c}_0 = \frac{1}{2} (\mu_0 - \gamma \Sigma_{sx} c_1)^\top (\gamma \Sigma)^{-1} (\mu_0 - \gamma \Sigma_{sx} c_1) - \frac{1}{2} c_1^\top \Omega c_1 + \frac{1}{2} \text{Tr}(c_2 \Sigma_x) \quad (97) \\ -\dot{c}_1 = (\mu - \gamma \Sigma_{sx} c_2)^\top (\gamma \Sigma)^{-1} (\mu - \gamma \Sigma_{sx} c_1) - c_2 \Omega c_1 - \kappa^\top c_1 \\ -\dot{c}_2 = (\mu - \gamma \Sigma_{sx} c_2)^\top (\gamma \Sigma)^{-1} (\mu - \gamma \Sigma_{sx} c_2) - c_2 \Omega c_2 - c_2 \kappa - \kappa^\top c_2$$

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 Riccati type:

$$-\dot{c}_1 = (\mu - \gamma \Sigma_{sx} c_2)^\top (\gamma \Sigma)^{-1} \mu_0 - \{(\mu - \gamma \Sigma_{sx} c_2)^\top \Sigma^{-1} \Sigma_{sx} + c_2^\top \Omega + \kappa^\top\} c_1 \quad (98) \\ -\dot{c}_2 = c_2^\top (\gamma \Sigma_{sx}^\top \Sigma^{-1} \Sigma_{sx} - \Omega) c_2 - c_2 (\kappa + \Sigma_{sx} \Sigma^{-1} \mu) - (\kappa + \Sigma_{sx} \Sigma^{-1} \mu)^\top c_2 + \mu^\top (\gamma \Sigma)^{-1} \mu \quad (99)$$

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

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

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 corollary 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:

$$CMVE_t = (\gamma \Sigma)^{-1} (\mu_0 + \mu x_t) \quad (100)$$

Remark 10 *The fact that the solution of the HJB equation (if it exists) solves the optimization problem at hand follows from a standard verification argument. Indeed, suppose there exists a*

²²Recall that for $x^\top C x = x^\top B x$ it is sufficient that $C = \frac{B+B^\top}{2}$, which insures that C is symmetric even if B is not. In fact, it is well-known that for a matrix A that is not symmetric, $C = \frac{A+A^\top}{2}$ is the unique matrix that is symmetric and such that $x^\top A x = x^\top C x \forall x$.

solution to the system of equation, from the definition of the HJB equation, it follows that

$$\mathbb{E}_t \left[dJ(x_t, t) + \{dW_t - \frac{1}{2}\gamma n_t^\top \Sigma n_t dt - \frac{1}{2}J_x^\top \Omega J_x dt - \gamma n_t^\top \Sigma_{sx} J_x dt\} \right] \leq 0 \quad \forall n_t$$

with equality for the optimal strategy. Integrating, this implies (given that $J(x, T) = 0 \quad \forall x$) that

$$J(x_t, t) \geq \mathbb{E} \left[\int_t^T \{dW_u - \frac{1}{2}\gamma n_u^\top \Sigma n_u du - \frac{1}{2}J_x^\top \Omega J_x du - \gamma n_u^\top \Sigma_{sx} J_x du\} \right] \quad \forall n_u \quad \text{and} \quad \forall t$$

and with equality for the optimal n_u^* .

This proves that our value function (together with its attached optimal strategy) is indeed dynamically optimal. However, given the recursive nature of the solution, it does not immediately follow from the verification theorem that there cannot be another (e.g., non-quadratic) function $J(x, t)$ (attached to a different optimal strategy), that might solve the same dynamic optimization problem and achieve a different (potentially higher) optimum (because the unknown J appears on both sides of the equation). Unlike in the classic optimization problems (such as the standard Merton portfolio problem), the uniqueness of the value function is not implied by the verification theorem. However, for the cases that are of interest to us (namely $\gamma_x = \gamma$ and $\gamma_x = 0$), then the uniqueness of the solution can easily be proved (for $\gamma_x = \gamma$ see footnote 19 and for $\gamma_x = 0$ the objective function is not recursive as shown in corollary 3).

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 satisfy:

$$J(n_t, x_t, t) = \max_{\theta} \mathbb{E}_t \left[\int_t^T \left\{ dW_u - \frac{1}{2}\gamma n_u^\top \Sigma n_u du - \frac{1}{2}J_x^\top \Omega J_x du - \gamma n_u^\top \Sigma_{sx} J_x du \right\} \right] \quad (101)$$

where we define:

$$\Omega = \gamma \sigma_{xs} \sigma_{xs}^\top + \gamma_x \sigma_x \sigma_x^\top \quad (102)$$

$$\Sigma_{sx} = \sigma_s \sigma_{xs}^\top \quad (103)$$

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

$$0 = \max_{\theta} \mathbb{E}_t [dW_t - \frac{1}{2}\gamma n_t^\top \Sigma n_t dt - \frac{1}{2}J_x^\top \Omega J_x dt - \gamma n_t^\top \Sigma_{sx} J_x dt + dJ(t, n_t, x_t)] \quad (104)$$

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

$$0 = \max_{\theta} \left\{ n^{\top}(\mu_0 + \mu x) - \frac{1}{2}\theta^{\top}\Lambda\theta - \frac{1}{2}\gamma n^{\top}\Sigma n - \frac{1}{2}J_x^{\top}\Omega J_x - \gamma n^{\top}\Sigma_{sx}J_x + J_t + J_n^{\top}\theta - J_x^{\top}\kappa x + \frac{1}{2}\text{Tr}(J_{xx}\Sigma_x) \right\} \quad (105)$$

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 gradient with respect to n , and J_t the partial derivative with respect to t .

The first order condition is²³

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

Plugging back into the HJB equation we get:

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

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 $Q(t)$ and $c_2(t)$ are symmetric (respectively N- and K-dimensional) matrices, $q(t)$ is a $(K \times N)$ matrix, $q_0(t)$ and $c_1(t)$ are vectors and $c_0(t)$ is a scalar (all of deterministic functions).

$$\begin{aligned} J_t &= -\frac{1}{2}n^{\top}\dot{Q}n + n^{\top}(\dot{q}_0 + \dot{q}^{\top}x) + \dot{c}_0 + \dot{c}_1^{\top}x + \frac{1}{2}x^{\top}\dot{c}_2x \\ J_n &= -Qn + q_0 + q^{\top}x \\ J_x &= qn + c_1 + c_2x \\ J_{xx} &= c_2 \end{aligned}$$

Thus HJB becomes:

$$\begin{aligned} 0 &= -\frac{1}{2}n^{\top}\dot{Q}n + n^{\top}(\dot{q}_0 + \dot{q}^{\top}x) + \dot{c}_0 + \dot{c}_1^{\top}x + \frac{1}{2}x^{\top}\dot{c}_2x \\ &\quad + \frac{1}{2}(-Qn + q_0 + q^{\top}x)^{\top}\Lambda^{-1}(-Qn + q_0 + q^{\top}x) + n^{\top}(\mu_0 + \mu x) - \frac{1}{2}\gamma n^{\top}\Sigma n \\ &\quad - \frac{1}{2}(qn + c_1 + c_2x)^{\top}\Omega(qn + c_1 + c_2x) - \gamma n^{\top}\Sigma_{sx}(qn + c_1 + c_2x) - (qn + c_1 + c_2x)^{\top}\kappa x + \frac{1}{2}\text{Tr}(c_2\Sigma_x) \end{aligned}$$

²³The second order condition is $\Lambda > 0$ and always satisfied.

Rewriting:

$$\begin{aligned}
0 = & \frac{1}{2}n^\top(-\dot{Q} + Q\Lambda^{-1}Q - \gamma\Sigma - q^\top\Omega q - 2\gamma\Sigma_{sx}q)n \\
& + n^\top(\dot{q}_0 + \mu_0 - Q\Lambda^{-1}q_0 - q^\top\Omega c_1 - \gamma\Sigma_{sx}c_1) + n^\top(\dot{q}^\top - Q\Lambda^{-1}q^\top + \mu - q^\top\kappa - q^\top\Omega c_2 - \gamma\Sigma_{sx}c_2)x \\
& + x^\top\left(\frac{1}{2}\dot{c}_2 + \frac{1}{2}q\Lambda^{-1}q^\top - c_2\kappa - \frac{1}{2}c_2\Omega c_2\right)x \\
& + x^\top(\dot{c}_1 + q\Lambda^{-1}q_0 - c_2\Omega c_1 - \kappa^\top c_1) + \dot{c}_0 + \frac{1}{2}q_0^\top\Lambda^{-1}q_0 - \frac{1}{2}c_1^\top\Omega c_1 + \frac{1}{2}\text{Tr}(c_2\Sigma_x)
\end{aligned}$$

So we obtain the set of ODEs that our candidate solution should satisfy:²⁴

$$-\dot{Q} = \gamma\Sigma - Q\Lambda^{-1}Q + q^\top\Omega q + \gamma(\Sigma_{sx}q + q^\top\Sigma_{sx}^\top) \quad (106)$$

$$-\dot{q}^\top = \mu - q^\top\kappa - Q\Lambda^{-1}q^\top - q^\top\Omega c_2 - \gamma\Sigma_{sx}c_2 \quad (107)$$

$$-\dot{c}_2 = -(c_2\kappa + \kappa^\top c_2) + q\Lambda^{-1}q^\top - c_2\Omega c_2 \quad (108)$$

$$-\dot{q}_0 = \mu_0 - Q\Lambda^{-1}q_0 - q^\top\Omega c_1 - \gamma\Sigma_{sx}c_1 \quad (109)$$

$$-\dot{c}_1 = -\kappa^\top c_1 + q\Lambda^{-1}q_0 - c_2\Omega c_1 \quad (110)$$

$$-\dot{c}_0 = \frac{1}{2}\text{Tr}(c_2\Sigma_x) + \frac{1}{2}q_0^\top\Lambda^{-1}q_0 - \frac{1}{2}c_1^\top\Omega c_1 \quad (111)$$

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, as 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:

$$dn_t = \tau_t(aim(x_t, t) - n_t)dt \quad (112)$$

²⁴See also footnote 22.

G The finite horizon solution with CMV preferences

As discussed in corollary (3), 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:

$$-\dot{Q} = \gamma\Sigma - Q\Lambda^{-1}Q \quad (113)$$

$$-\dot{q}^\top = \mu - q^\top\kappa - Q\Lambda^{-1}q^\top \quad (114)$$

$$-\dot{q}_0 = \mu_0 - Q\Lambda^{-1}q_0 \quad (115)$$

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

$$-\dot{\tau} = \gamma\Lambda^{-1}\Sigma - \tau\tau \quad (116)$$

$$-\dot{Q} = \gamma\Sigma - \tau^\top Q \quad (117)$$

$$-\dot{q}^\top = \mu - q^\top\kappa - \tau^\top q^\top \quad (118)$$

$$-\dot{q}_0 = \mu_0 - \tau^\top q_0 \quad (119)$$

This system has an intuitive closed-form solution in terms of the eigenvector, eigenvalue decomposition of the matrix $\gamma\Lambda^{-1}\Sigma = FD_\eta F^{-1}$ defined in equation (27). Indeed, if we define $\tau_t = FD_{h_t}F^{-1}$ where D_{h_t} 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 $h_i(t)$ functions, which solve:

$$-\dot{h}_i = \eta_i - h_i^2 \quad s.t. \quad h_i(T) = 0 \quad (120)$$

The solution is then as given in the theorem. It follows that the trading speed matrix is given by $\tau_t = FD_{h_t}(t)F^{-1}$ and the Q matrix is $Q(t) = \Lambda\tau(t)$.

To solve for $q(t), q_0(t)$, we use the following lemmas.

lemma 1 *Since τ_t is diagonalizable then so is its inverse:*

$$\tau_t^\top = \tilde{F}D_{h_t}\tilde{F}^{-1} \quad (121)$$

where $\tilde{F} = F^{-\top}$ is the inverse of the transpose of F (or equivalently the transpose of the inverse of F).

Proof. From its solution $\tau_t^\top = (F^{-1})^\top D_{h_t} F^\top$. Thus $\tilde{F} = (F^{-1})^\top$. It remains to show that $(F^{-1})^\top = (F^\top)^{-1}$ for then $\tilde{F}^{-1} = F^\top$ and the decomposition obtains. But note that $(F^{-1})^\top F^\top = (FF^{-1})^\top = I$.

Clearly, if Σ and Λ have same eigenfactors or if either is a diagonal matrix, then $F = \tilde{F}$. But in general this need not be the case, as $\gamma\Lambda^{-1}\Sigma$ (and therefore τ_t) need not be symmetric. ■

lemma 2 *The following holds:*

$$\frac{d}{dt}e^{-\int_0^t \tau_s^\top ds} = -\tau_t^\top e^{-\int_0^t \tau_s^\top ds} dt = -e^{-\int_0^t \tau_s^\top ds} \tau_t^\top dt \quad (122)$$

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

$$e^{-\int_t^u \tau_s^\top ds} e^{-\int_u^T \tau_s^\top ds} = e^{-\int_t^T \tau_s^\top ds} \quad (123)$$

Proof. Note that $\int_0^t \tau_s^\top ds = \tilde{F}D_{\int_0^t h(s)ds} \tilde{F}^{-1}$. Therefore, from the properties of the matrix exponential²⁵ $e^{-\int_0^t \tau_s^\top ds} = \tilde{F}D_{e^{-\int_0^t h(s)ds}} \tilde{F}^{-1}$. Now, taking the derivative we find:

$$\frac{d}{dt}e^{-\int_0^t \tau_s^\top ds} = \tilde{F}D_{e^{-\int_0^t h(s)ds}h(t)} \tilde{F}^{-1} \quad (124)$$

$$= \tilde{F}D_{e^{-\int_0^t h(s)ds}} \tilde{F}^{-1} \tilde{F}D_{-h(t)} \tilde{F}^{-1} \quad (125)$$

$$= e^{-\int_0^t \tau_s^\top ds} \tau_t \quad (126)$$

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:

$$e^{-\int_t^u \tau_s^\top ds} e^{-\int_u^T \tau_s^\top ds} = \tilde{F}D_{e^{-\int_t^u h_s ds}} \tilde{F}^{-1} \tilde{F}D_{e^{-\int_u^T h_s ds}} \tilde{F}^{-1} \quad (127)$$

$$= \tilde{F}D_{e^{-\int_t^T h_s ds}} \tilde{F}^{-1} \quad (128)$$

$$= e^{-\int_t^T \tau_s^\top ds} \quad (129)$$

■

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

²⁵The 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 $e^{YXY^{-1}} = Y e^X Y^{-1}$.

lemma 3 *The solution to the ODE system is given as follows;*

$$q(t) = \int_t^T e^{-\int_t^u \tau_s^\top ds} \mu e^{-\int_t^u \kappa ds} du \quad (130)$$

$$q_0(t) = \int_t^T e^{-\int_t^u \tau_s^\top ds} du \mu_0 \quad (131)$$

$$Q(t) = \int_t^T e^{-\int_t^u \tau_s^\top ds} du \gamma \Sigma \quad (132)$$

Proof. We prove only the first results as the other ones are proved similarly. Using lemma 2 we have:

$$\frac{d}{dt} e^{\int_0^t \tau_s^\top ds} q(t)^\top e^{\int_0^t \kappa ds} = e^{\int_0^t \tau_s^\top ds} (\tau_t^\top q(t)^\top + q(t)^\top \dot{\cdot} + q(t)^\top \kappa) e^{\int_0^t \kappa ds} \quad (133)$$

$$= e^{\int_0^t \tau_s^\top ds} \mu e^{\int_0^t \kappa ds} \quad (134)$$

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

$$e^{\int_0^t \tau_s^\top ds} q(t)^\top e^{\int_0^t \kappa ds} = - \int_t^T e^{\int_0^u \tau_s^\top ds} \mu e^{\int_0^u \kappa ds} du \quad (135)$$

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 2 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

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:

$$J(x_t) = \max_n E_t \left[\int_t^\infty e^{-\rho(u-t)} \left\{ dW_u - \frac{1}{2} \gamma n_u^\top \Sigma n_u du - \frac{1}{2} J_x^\top \Omega J_x du - \gamma n_u^\top \Sigma_{sx} J_x du \right\} \right] \quad (136)$$

The corresponding Bellman-equation is:

$$0 = \max_n E_t [dW_t - \frac{1}{2} \gamma n_t^\top \Sigma n_t dt - \frac{1}{2} J_x^\top \Omega J_x dt - \gamma n_t^\top \Sigma_{sx} J_x dt + dJ(x_t) - \rho J(x_t)] \quad (137)$$

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

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

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 , is²⁶

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

Plugging back into the HJB equation we get:

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

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 a symmetric positive definite matrix, c_1 is a K -dimensional vector, and c_0 a constant.

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

Thus the HJB equation becomes

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

This equation is satisfied if c_0, c_1, c_2 solve the following system:²⁷

$$\begin{aligned} \rho c_0 &= \frac{1}{2} (\mu_0 - \gamma \Sigma_{sx} c_1)^\top (\gamma \Sigma)^{-1} (\mu_0 - \gamma \Sigma_{sx} c_1) - \frac{1}{2} c_1^\top \Omega c_1 + \frac{1}{2} \text{Tr}(c_2 \Sigma_x) \\ \rho c_1 &= (\mu - \gamma \Sigma_{sx} c_2)^\top (\gamma \Sigma)^{-1} (\mu_0 - \gamma \Sigma_{sx} c_1) - c_2 \Omega c_1 - \kappa^\top c_1 \\ \rho c_2 &= (\mu - \gamma \Sigma_{sx} c_2)^\top (\gamma \Sigma)^{-1} (\mu - \gamma \Sigma_{sx} c_2) - c_2 \Omega c_2 - c_2 \kappa - \kappa^\top c_2 \end{aligned} \quad (138)$$

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 Riccati type:

$$0 = (\mu - \gamma \Sigma_{sx} c_2)^\top (\gamma \Sigma)^{-1} \mu_0 - \{(\mu - \gamma \Sigma_{sx} c_2)^\top \Sigma^{-1} \Sigma_{sx} + c_2^\top \Omega + \kappa^\top + \rho\} c_1 \quad (139)$$

$$0 = c_2 \left(\gamma \Sigma_{sx}^\top \Sigma^{-1} \Sigma_{sx} - \Omega \right) c_2 - c_2 (\rho + \kappa + \Sigma_{sx} \Sigma^{-1} \mu) - (\kappa + \Sigma_{sx} \Sigma^{-1} \mu)^\top c_2 + \mu^\top (\gamma \Sigma)^{-1} \mu \quad (140)$$

²⁶The second order condition: $\gamma \Sigma > 0$ is always satisfied.

²⁷See also footnote 22.

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

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

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 corollary 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:

$$Mwz_t = (\gamma \Sigma)^{-1}(\mu_0 + \mu x_t) \quad (141)$$

Remark 11 *To apply a standard verification theorem we require in addition that the transversality condition $\lim_{T \rightarrow \infty} \mathbb{E}[e^{-\rho T} J(X_T)] = 0$ be satisfied. 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*

$$\mathbb{E}_t \left[de^{-\rho t} J(x_t) + e^{-\rho t} \left\{ dW_t - \frac{1}{2} \gamma n_t^\top \Sigma n_t dt - \frac{1}{2} J_x^\top \Omega J_x dt - \gamma n_t^\top \Sigma_{sx} J_x dt \right\} \right] \leq 0 \quad \forall n_t$$

with equality at the optimal strategy. This implies that

$$J(x_t) \geq \mathbb{E}_t[e^{-\rho T} J(x_T)] + \mathbb{E} \left[\int_t^T e^{-\rho u} \left\{ dW_u - \frac{1}{2} \gamma n_u^\top \Sigma n_u du - \frac{1}{2} J_x^\top \Omega J_x du - \gamma n_u^\top \Sigma_{sx} J_x du \right\} \right] \quad \forall n_u \text{ and } \forall t$$

and with equality for the optimal n_u . Letting $T \rightarrow \infty$ using the transversality condition establishes the dynamic optimality of the value function (and of the of the associated trading strategy). As discussed in the final horizon case (see remark 10) there is an issue regarding the uniqueness of the value function. In fact, there appear to be several solutions to the system of (quadratic) equations derived from the HJB equation. We focus on the solution (selecting the positive definite Q and c_2 matrices) that is most economically sensible and consistent with the finite horizon solution. We conjecture (but could not prove for the case where $\gamma_x \neq 0$) that, as in the finite horizon case, this is the unique solution consistent with some appropriately defined transversality condition (e.g., resulting from the finiteness of the expected utility).

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

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

$$J(n_t, x_t) = \max_{\theta} \mathbb{E}_t \left[\int_t^{\infty} e^{-\rho(s-t)} \left\{ dW_u - \frac{1}{2} \gamma n_u^\top \Sigma n_u du - \frac{1}{2} J_x^\top \Omega J_x du - \gamma n_u^\top \Sigma_{sx} J_x du \right\} \right] \quad (142)$$

where we define:

$$\Omega = \gamma \sigma_{xs} \sigma_{xs}^\top + \gamma_x \sigma_x \sigma_x^\top \quad (143)$$

$$\Sigma_{sx} = \sigma_s \sigma_{xs}^\top \quad (144)$$

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

$$0 = \max_{\theta} \mathbb{E}_t [dW_t - \frac{1}{2} \gamma n_t^\top \Sigma n_t dt - \frac{1}{2} J_x^\top \Omega J_x dt - \gamma n_t^\top \Sigma_{sx} J_x dt + dJ(n_t, x_t) - \rho J(n_t, x_t) dt] \quad (145)$$

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

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

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 gradient with respect to n .

The first order condition is:²⁸

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

Plugging back into the HJB equation we get:

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

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 and c_2 are symmetric positive-definite (respectively N- and K-dimensional) matrices,

²⁸The second order condition $\Lambda > 0$ is always satisfied.

q is a $(K \times N)$ matrix, q_0 and c_1 are vectors and c_0 is a scalar (all constant).

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

Thus HJB becomes

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

Rewriting:

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

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

$$\rho Q = \gamma \Sigma - Q \Lambda^{-1} Q + q^\top \Omega q + \gamma (\Sigma_{sx} q + q^\top \Sigma_{sx}^\top) \quad (146)$$

$$\rho q^\top = \mu - q^\top \kappa - Q \Lambda^{-1} q^\top - q^\top \Omega c_2 - \gamma \Sigma_{sx} c_2 \quad (147)$$

$$\rho c_2 = -(c_2 \kappa + \kappa^\top c_2) + q \Lambda^{-1} q^\top - c_2 \Omega c_2 \quad (148)$$

$$\rho c_0 = \frac{1}{2} \text{Tr}(c_2 \Sigma_x) + \frac{1}{2} q_0^\top \Lambda^{-1} q_0 - \frac{1}{2} c_1^\top \Omega c_1 \quad (149)$$

$$\rho q_0 = \mu_0 - Q \Lambda^{-1} q_0 - q^\top \Omega c_1 - \gamma \Sigma_{sx} c_1 \quad (150)$$

$$\rho c_1 = -\kappa^\top c_1 + q \Lambda^{-1} q_0 - c_2 \Omega c_1 \quad (151)$$

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:

$$dn_t = \tau(aim(x_t) - n_t)dt \quad (152)$$

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:

$$0 = c_2 \left(\gamma \Sigma_{sx}^\top \Sigma^{-1} \Sigma_{sx} - \Omega \right) c_2 - c_2 (\rho + 2\kappa + 2\Sigma_{sx} \Sigma^{-1} \mu) + \mu^\top (\gamma \Sigma)^{-1} \mu \quad (153)$$

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:

$$0 = c_2^2 \gamma_x \sigma_x^2 - c_2 (\rho + 2\kappa + 2 \frac{\sigma_{xs}}{\sigma_s} \mu) + \frac{\mu^2}{\gamma \sigma_s^2} \quad (154)$$

The positive 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:

$$\rho Q = \gamma \Sigma - Q \Lambda^{-1} Q + q^\top \Omega q + 2\gamma \Sigma_{sx} q \quad (155)$$

$$\rho q^\top = \mu - q^\top \kappa - Q \Lambda^{-1} q^\top - q^\top \Omega c_2 - \gamma \Sigma_{sx} c_2 \quad (156)$$

$$\rho c_2 = -2c_2 \kappa + q \Lambda^{-1} q^\top - c_2 \Omega c_2 \quad (157)$$

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:

$$\rho Q = \gamma\Sigma - Q\Lambda^{-1}Q + q^2\Omega + 2\gamma\Sigma_{sx}q \quad (158)$$

$$0 = \mu - q(\kappa + \rho) - Q\Lambda^{-1}q - q\Omega c_2 - \gamma\Sigma_{sx}c_2 \quad (159)$$

$$0 = -(\rho + 2\kappa)c_2 + q^2\Lambda^{-1} - c_2^2\Omega \quad (160)$$

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

$$\tau^2 + \rho\tau = \gamma\Lambda^{-1}\Sigma + q^2\Lambda^{-1}\Omega + 2\gamma\Lambda^{-1}\Sigma_{sx}q \quad (161)$$

$$0 = \mu - q(\kappa + \rho) - \tau q - q\Omega c_2 - \gamma\Sigma_{sx}c_2 \quad (162)$$

$$0 = (\rho + 2\kappa)c_2 - q^2\Lambda^{-1} + c_2^2\Omega \quad (163)$$

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:

$$\tau(q) = \frac{-\rho + \sqrt{\rho^2 + 4\Lambda^{-1}(q^2\gamma_x\sigma_x^2 + \gamma(\sigma_s + q\sigma_{xs})^2)}}{2} \quad (164)$$

$$c_2(q) = \frac{-(\rho + 2\kappa) + \sqrt{(\rho + 2\kappa)^2 + 4\Lambda^{-1}q^2(\gamma\sigma_{xs}^2 + \gamma_x\sigma_x^2)}}{2\Omega} \quad (165)$$

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

$$q(\kappa + \rho) + \tau(q)q + (q\gamma_x\sigma_x^2 + \gamma\sigma_{xs}(q\sigma_{xs} + \sigma_s))c_2(q) = \mu \quad (166)$$

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 $q \rightarrow \infty$).

The optimal aim portfolio is given by

$$aim(x) = Q^{-1}qx \quad (167)$$

$$= \tau^{-1}\Lambda^{-1}qx \quad (168)$$

$$(169)$$

K The infinite horizon solution CMV preferences

As discussed in corollary 3, 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:

$$\rho Q = \gamma \Sigma - Q \Lambda^{-1} Q \quad (170)$$

$$\rho q^\top = \mu - q^\top \kappa - Q \Lambda^{-1} q^\top \quad (171)$$

$$\rho q_0 = \mu_0 - Q \Lambda^{-1} q_0 \quad (172)$$

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

$$\rho \tau = \gamma \Lambda^{-1} \Sigma - \tau \tau \quad (173)$$

$$\rho Q = \gamma \Sigma - \tau^\top Q \quad (174)$$

$$\rho q^\top = \mu - q^\top \kappa - \tau^\top q^\top \quad (175)$$

$$\rho q_0 = \mu_0 - \tau^\top q_0 \quad (176)$$

This system has an intuitive closed-form solution in term of the diagonalization of the matrix $\gamma \Lambda^{-1} \Sigma = F D_\eta F^{-1}$ defined in equation (27). Plugging into the system of equation we find that $\tau = F D_h F^{-1}$ where the $h_i \forall i = 1, \dots, N$ are the positive roots of the quadratic equations:

$$\rho h_i = \eta_i - h_i^2 \quad (177)$$

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 $\tau = F D_h F^{-1}$ 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 4 for the expression for q

$$Q = (\rho + \tau^\top)^{-1} \gamma \Sigma \quad (178)$$

$$q_0 = (\rho + \tau^\top)^{-1} \mu_0 \quad (179)$$

$$q^\top = \int_0^\infty e^{-(\rho + \tau^\top)t} \mu e^{-\kappa t} dt \quad (180)$$

lemma 4 Suppose A, B are (full rank) square matrices with strictly positive eigenvalues. Then the matrix equation $-AX - XB = -C$ has the solution $X = \int_0^\infty e^{-At} C e^{-Bt} dt$.

Proof. Note that:

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

Taking the limit as $T \rightarrow \infty$ and noting that, since all the eigenvalues of A, B are positive we have $\lim_{T \rightarrow \infty} e^{-AT} C e^{-BT} = 0$, 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 $aim(x) = Q^{-1}(q_0 + qx)$.