Online Appendix for:

Liquidity regimes and optimal dynamic asset allocation

Pierre Collin-Dufresne[a,*] Kent Daniel[b,c], and Mehmet Sağlam[d]

[a]Swiss Finance Institute at EPFL, Lausanne, Switzerland
[b]Columbia Business School, New York, NY, USA
[c]National Bureau of Economic Research, Cambridge, MA, USA
[d]University of Cincinnati, Cincinnati, OH, USA

Appendix A. Additional Figures and Tables

This section of the online appendix contains additional tables and figures, as referenced in the main text of the paper.

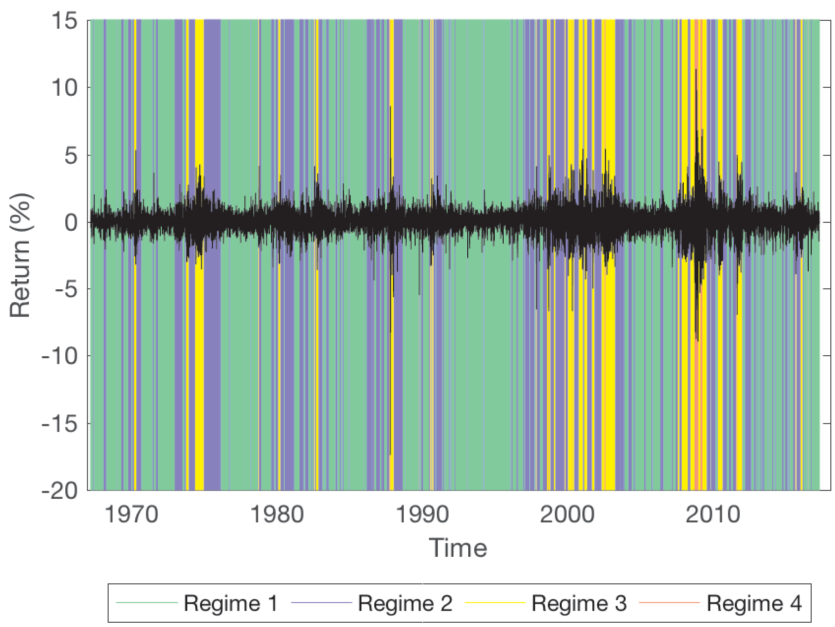

Fig. A.1. The colored vertical bars in this correspond to the regime with the maximum smoothed probability at that time. The state probabilities are plotted in figure . Daily excess market returns are plotted in black. The time period is 1967-03-13 to 2017-03-31.

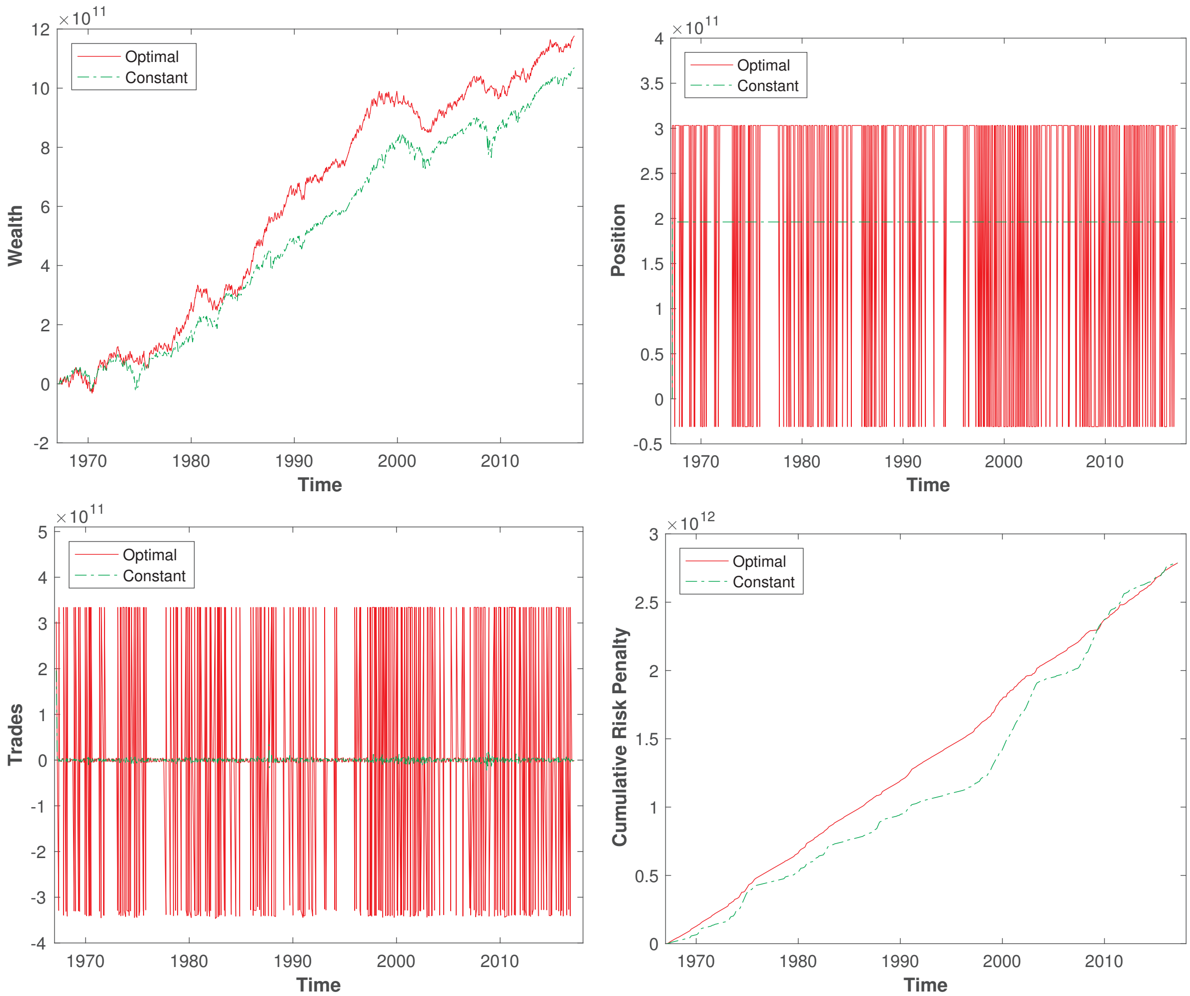

Fig. A.2. This figure compares the out-of-sample performance of the optimal policy with a constant-dollar portfolio in the absence of trading costs. Both strategies start from zero-wealth. Constant portfolio is equal to the scaled unconditional Markowitz portfolio so that the policy has the same risk exposure as the optimal policy (as shown in the bottom-right panel). Both strategies are constructed using only backward-looking data. Top-left panel shows the cumulative wealth of each policy which is the main comparison metric. Top-right panel shows the each strategy's dollar position in the market fund. Bottom-left panel illustrates the change in position due to the rebalancing of the strategy.

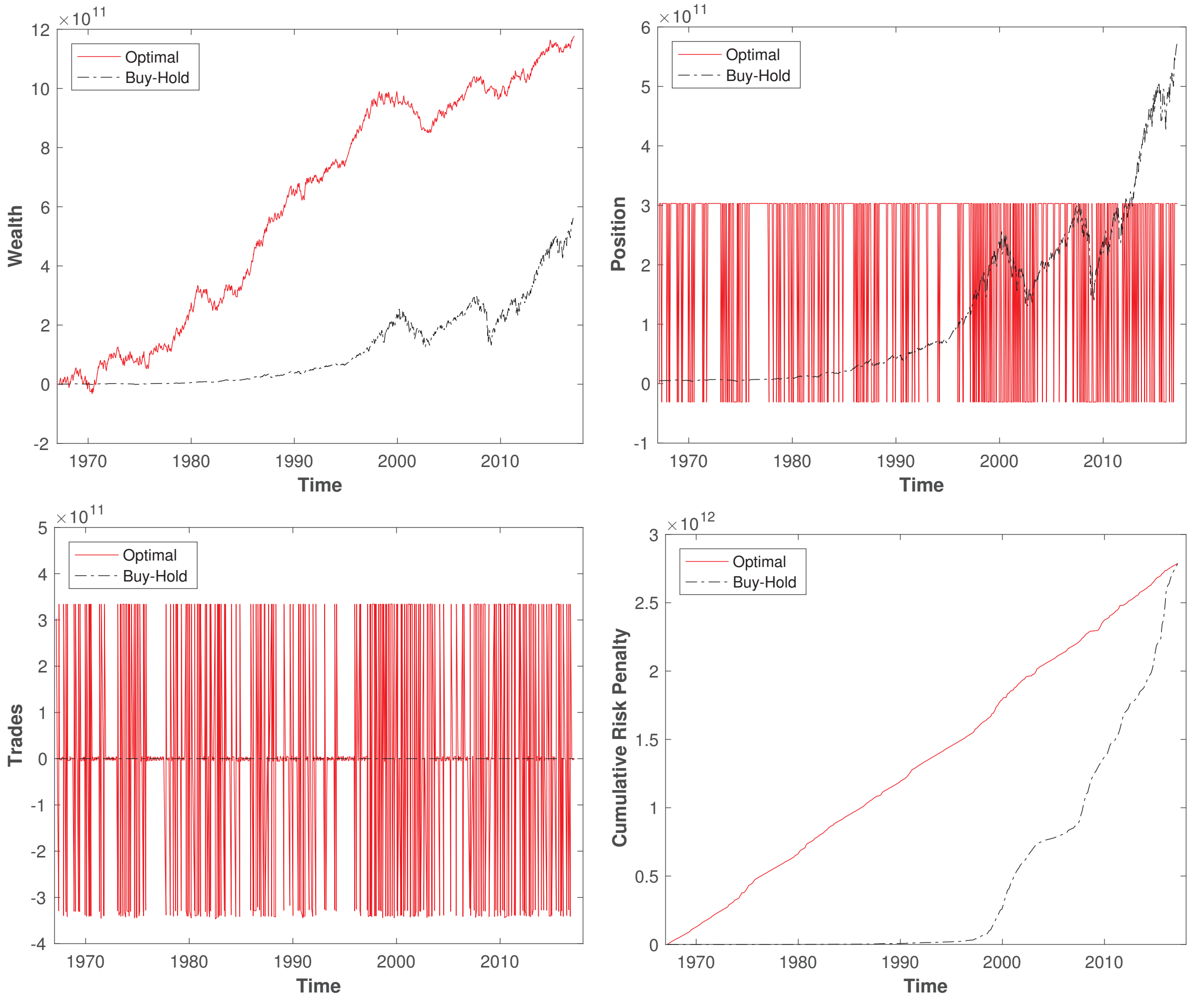

Fig. A.3. This figure compares the out-of-sample performance of the optimal policy with a buy-and-hold portfolio portfolio in the absence of trading costs. Both strategies start from zero-wealth. Buy-and-hold portfolio invests x_0 dollars into the market fund at the beginning of the horizon. We scale x_0 so that the policy has the same risk exposure as the optimal policy (as shown in the bottom-right panel). Both strategies are constructed using only backward-looking data. Top-left panel shows the cumulative wealth of each policy which is the main comparison metric. Top-right panel shows the each strategy's dollar position in the market fund. Bottom-left panel illustrates the change in position due to the rebalancing of the strategy.

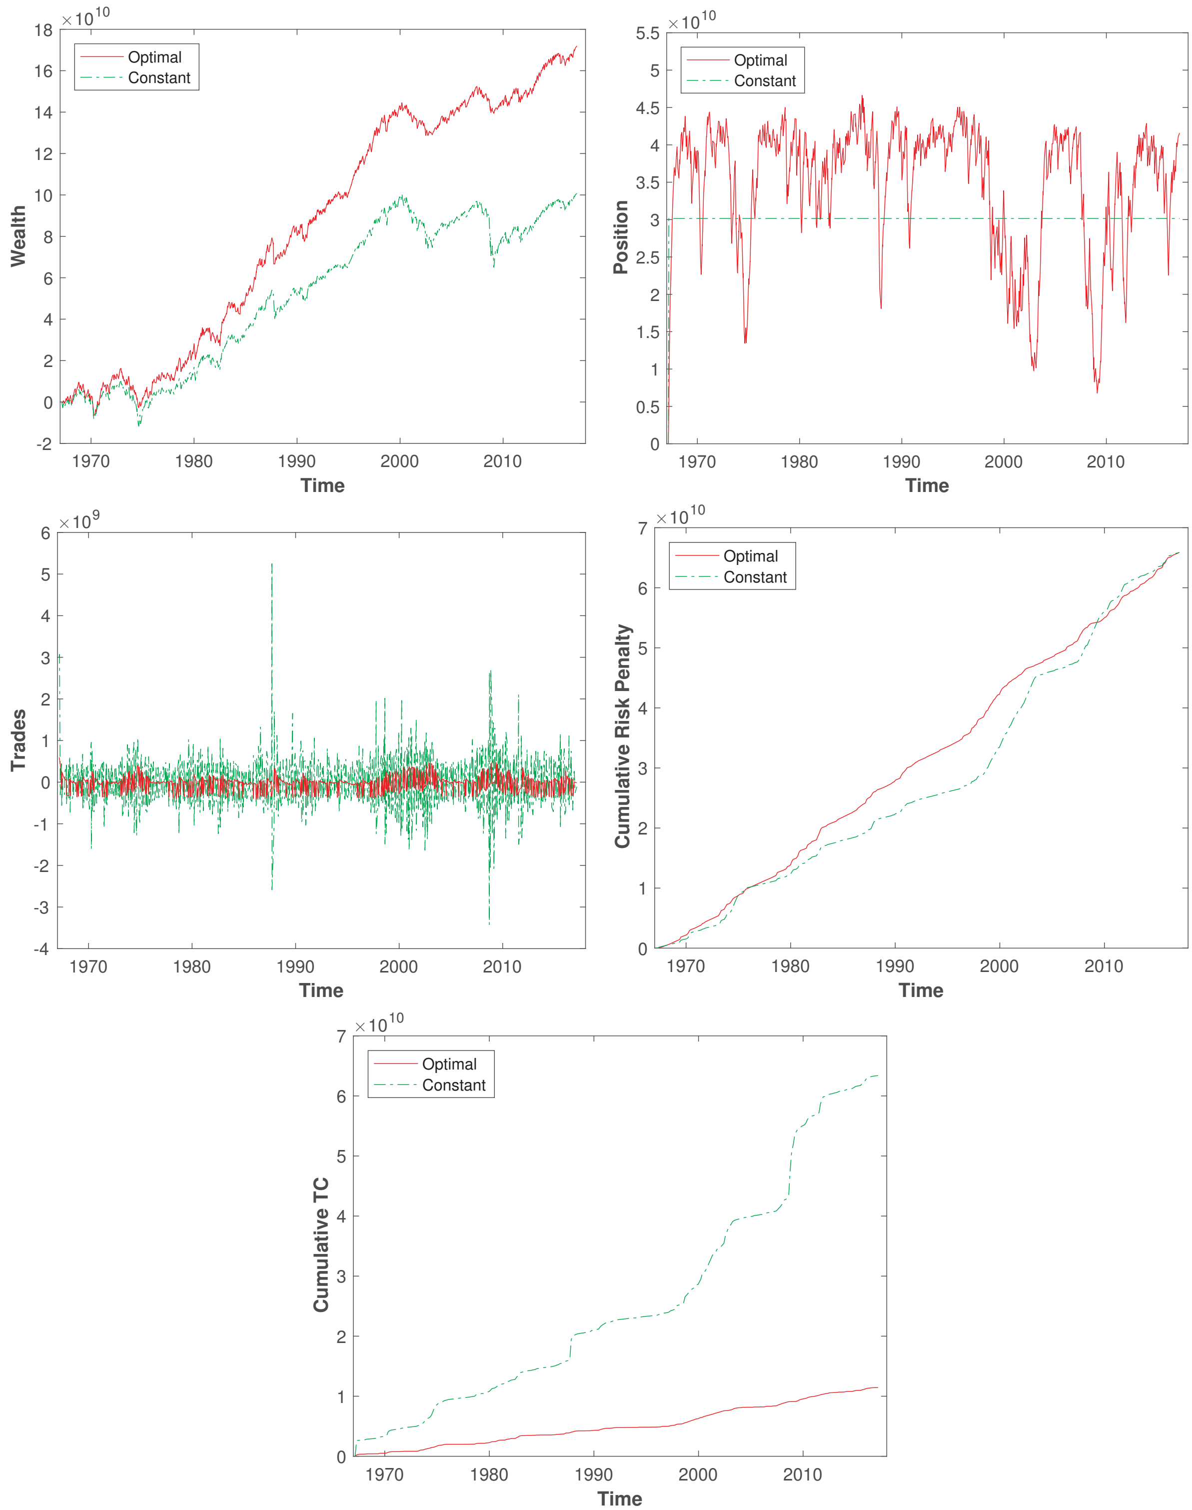

Fig. A.4. This figure compares the out-of-sample performance of the optimal policy with a constant-dollar portfolio in the presence of trading costs. Both strategies start from zero-wealth. Constant portfolio is equal to the scaled unconditional Markowitz portfolio so that the policy has the same risk exposure as the optimal policy (as shown in the center-right panel). Both strategies are constructed using only backward-looking data. Top-left panel shows the cumulative wealth of each policy which is the main comparison metric. Top-right panel shows the each strategy's dollar position in the market fund. Center-left panel illustrates the change in position due to the rebalancing of the strategy and the bottom panel illustrates the cumulative cost of these trades.

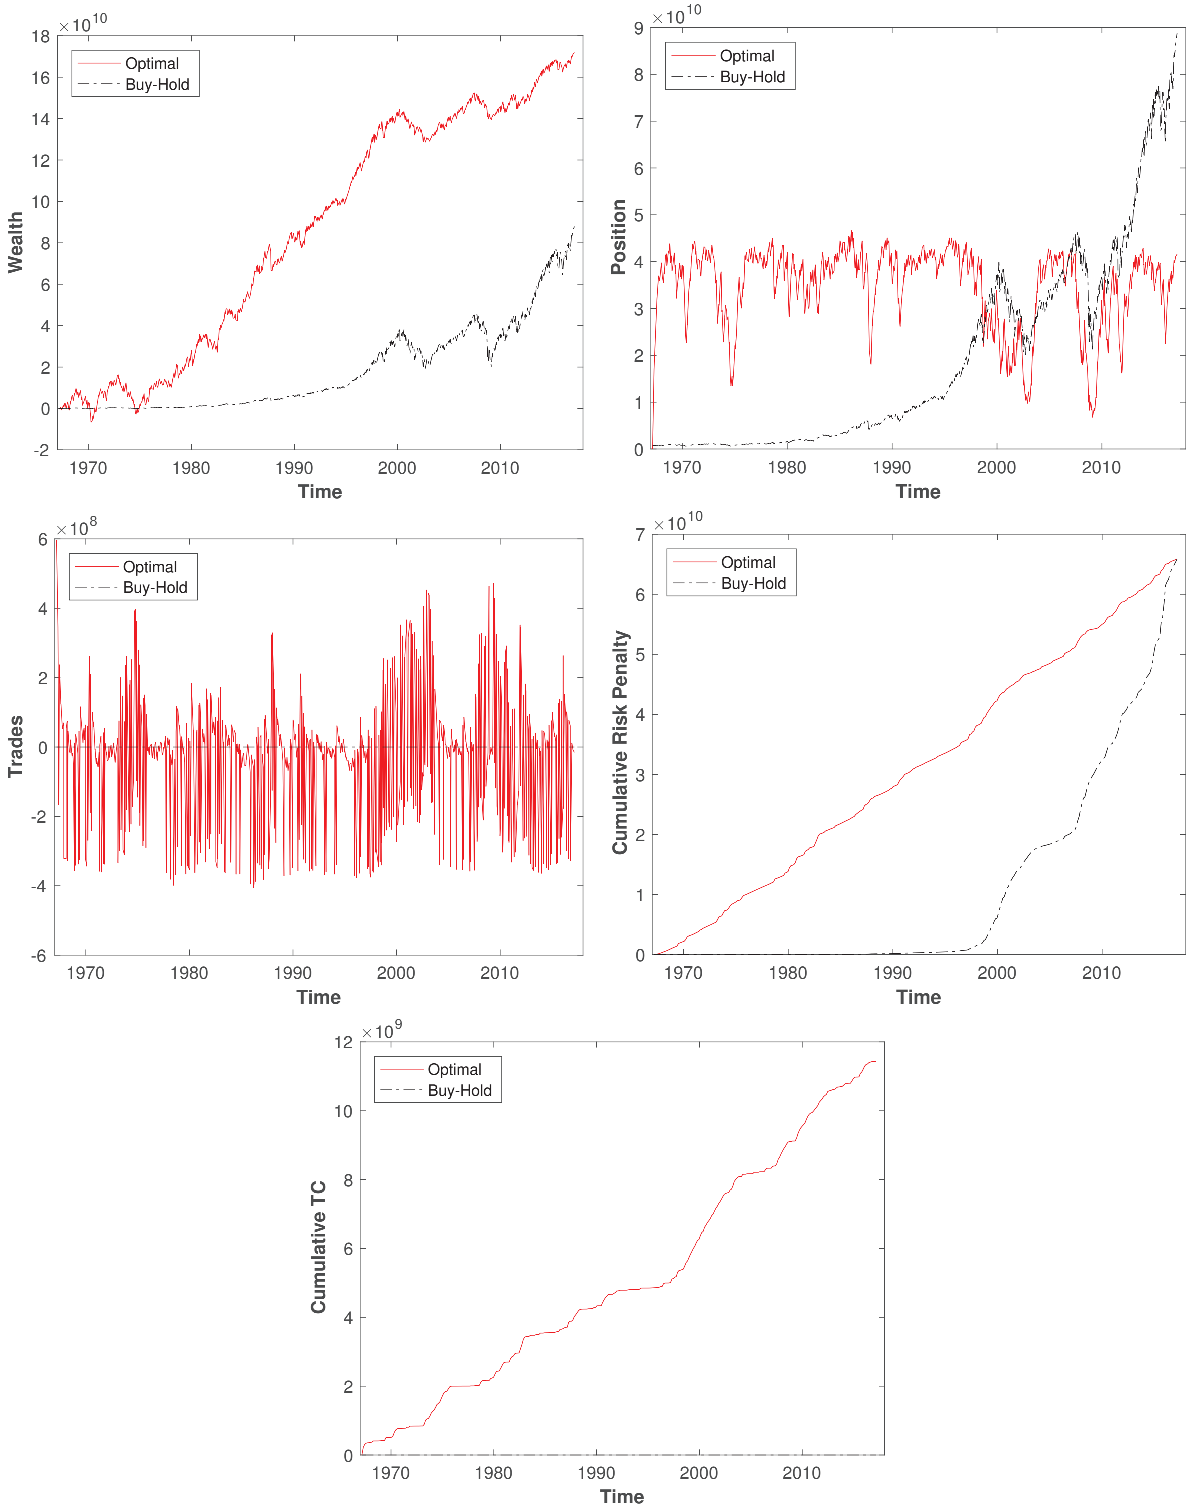

Fig. A.5. This figure compares the out-of-sample performance of the optimal policy with a buy-and-hold portfolio in the presence of trading costs. Both strategies start from zero-wealth. Buy-and-hold portfolio invests x_0 dollars into the market fund at the beginning of the horizon. We scale x_0 so that the policy has the same risk exposure as the optimal policy (as shown in the center-right panel). Both strategies are constructed using only backward-looking data. Top-left panel shows the cumulative wealth of each policy which is the main comparison metric. Top-right panel shows the each strategy's dollar position in the market fund. Center-left panel illustrates the change in position due to the rebalancing of the strategy and the bottom panel illustrates the cumulative cost of these trades.

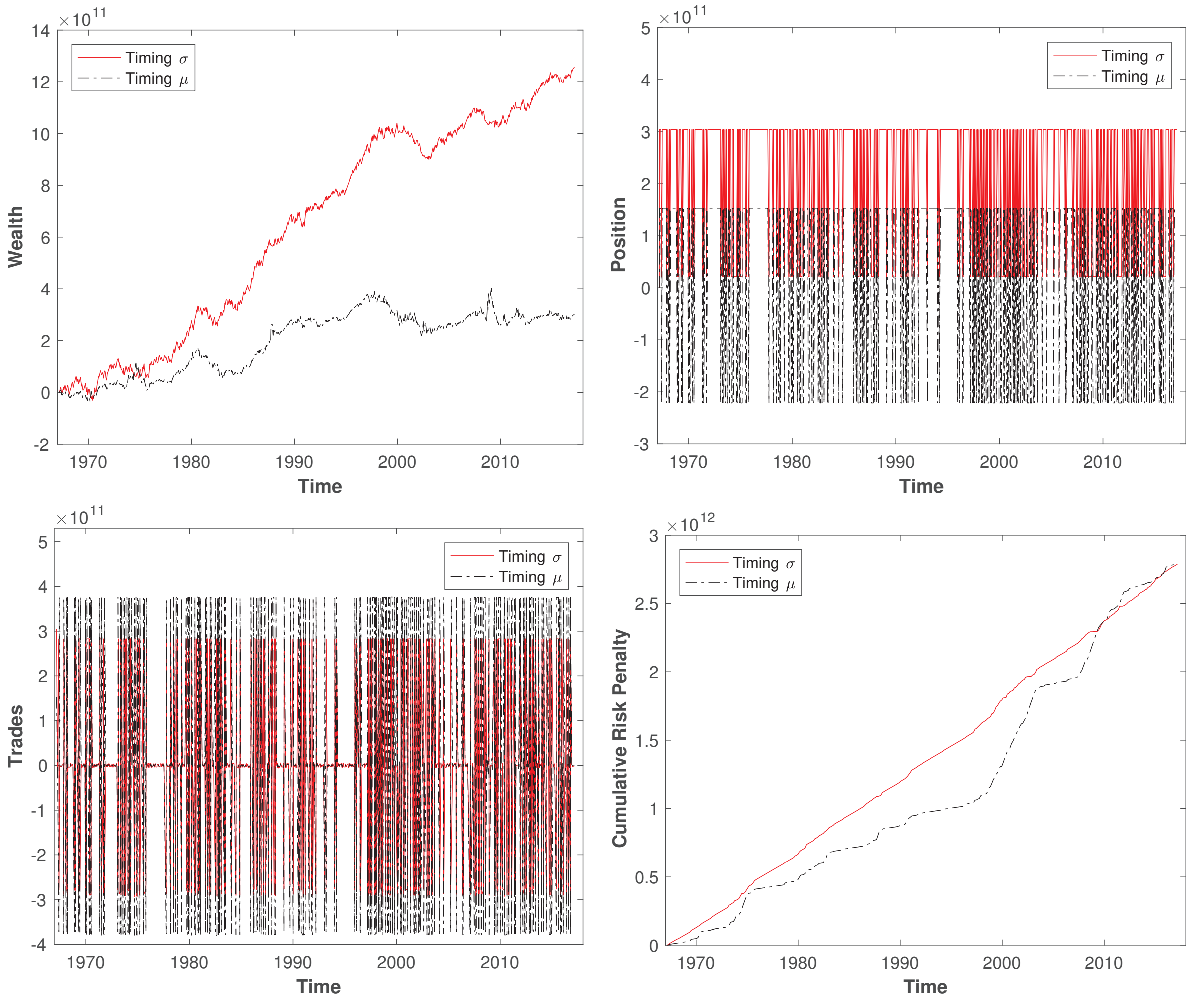

Fig. A.6. This figure compares the out-of-sample performance of timing strategies in the absence of trading costs. σ-timing policy takes into account that the volatility is time-varying between two states but assumes that expected return is constant throughout the investment horizon. μ-timing policy internalizes the potential switches in the expected returns but it assumes that the volatility stays constant at an unconditional average. Both strategies start from zero-wealth. We scale both policies so that they have the same risk exposure as the optimal policy (as shown in the bottom-right panel). Top-left panel shows the cumulative wealth of each policy which is the main comparison metric. Top-right panel shows the each strategy's dollar position in the market fund. Bottom-left panel illustrates the change in position due to the rebalancing of the strategy.

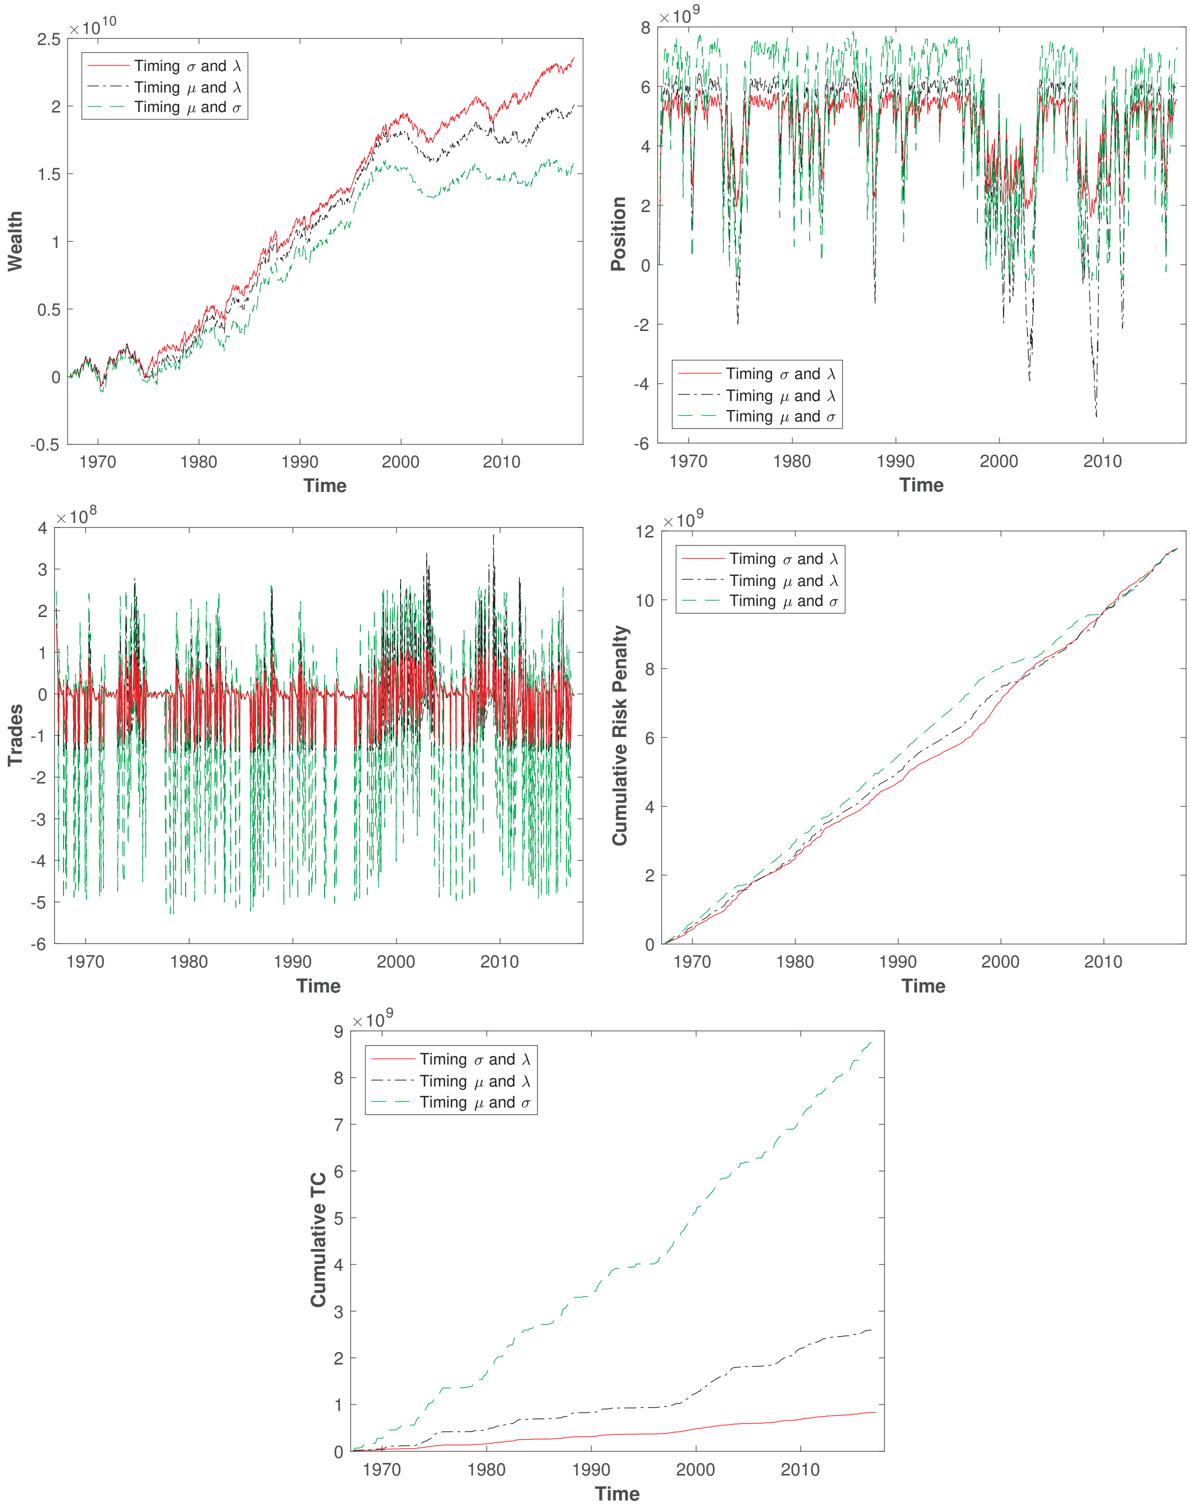

Fig. A.7. This figure compares the out-of-sample performance of timing strategies in the presence of trading costs. Timing σ and λ policy takes into account that the volatility and trading costs are time-varying between two states but assumes that expected return is constant at its unconditional average. Timing μ and λ policy internalizes the potential switches in the expected returns and transaction costs but it assumes that the volatility stays constant at an unconditional average. Finally, timing σ and μ policy takes into account that the volatility and expected returns are time-varying between two states but assumes that trading costs are constant at its unconditional average. Risk aversion level is at 1×10^{-9}.

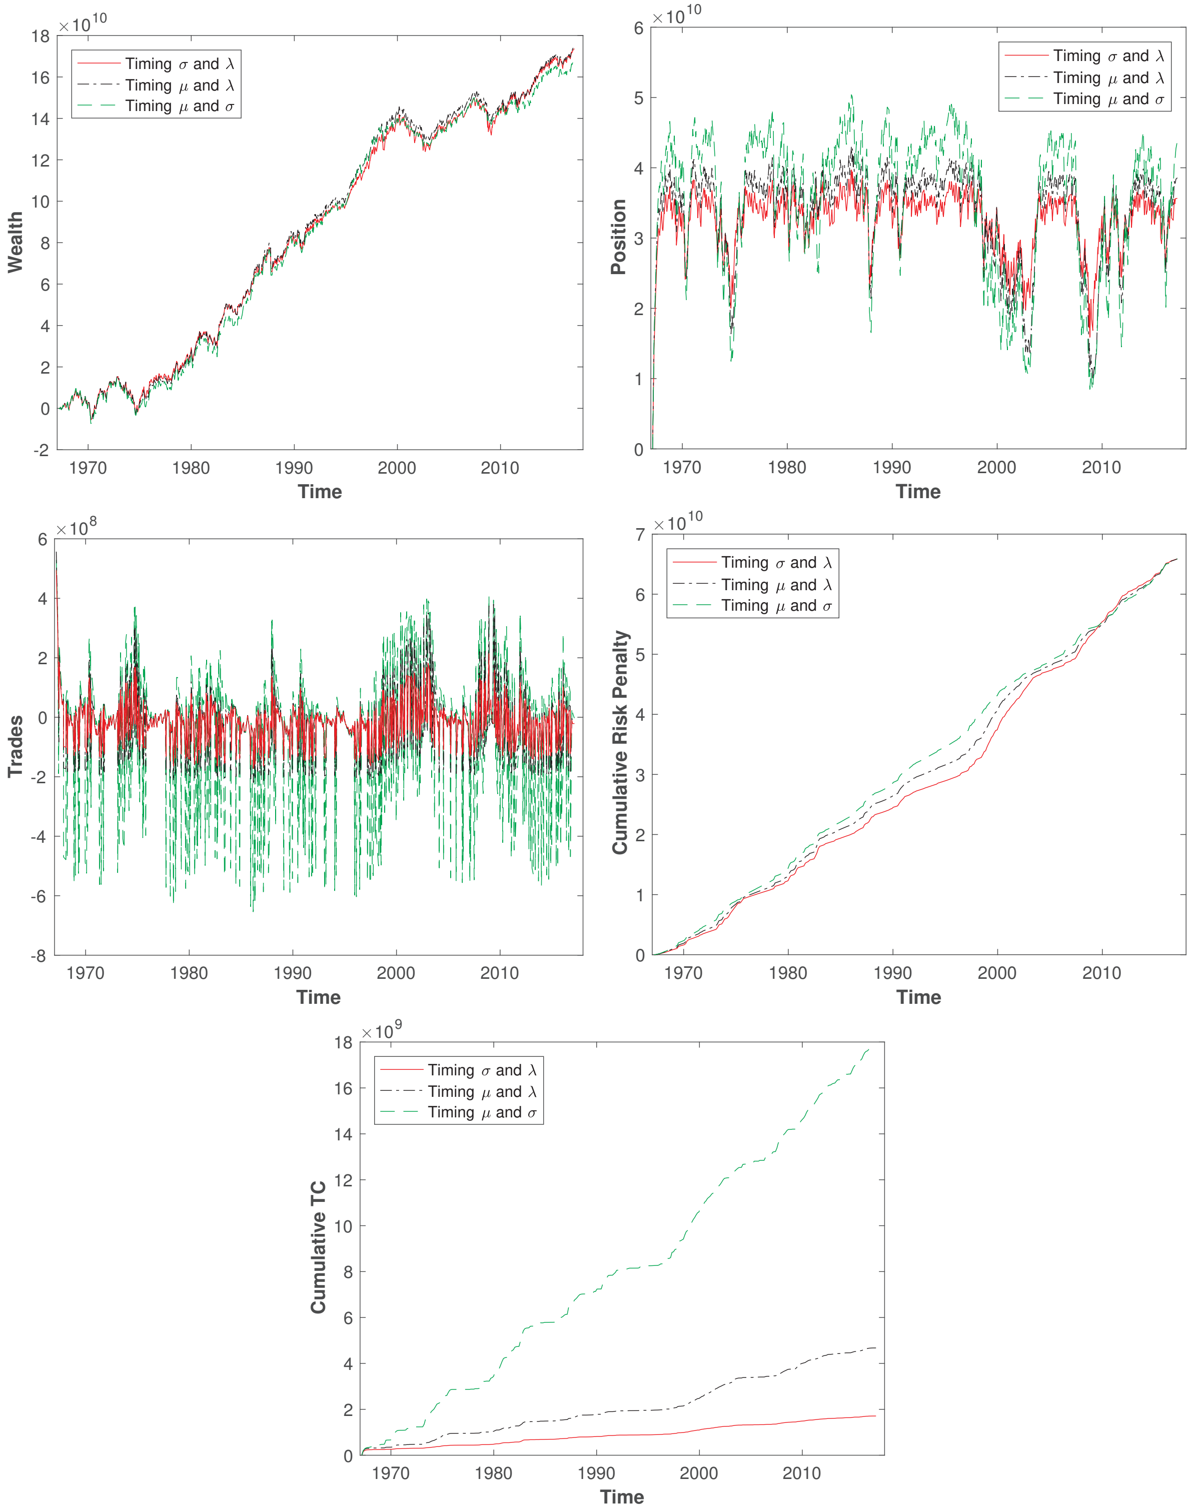

Fig. A.8. This figure compares the out-of-sample performance of timing strategies in the presence of trading costs. Timing σ and λ policy takes into account that the volatility and trading costs are time-varying between two states but assumes that expected return is constant at its unconditional average. Timing μ and λ policy internalizes the potential switches in the expected returns and transaction costs but it assumes that the volatility stays constant at an unconditional average. Finally, timing σ and μ policy takes into account that the volatility and expected returns are time-varying between two states but assumes that trading costs are constant at its unconditional average. Risk aversion level is at 1×10^{-10}.

Appendix B. Additional Derivations

We begin by solving a continuous-time version of the model presented in Section 2 and show that the model implications and insights are consistent with its discrete time counterpart. Finite- and infinite-horizon versions of the continuous-time model are presented in Appendices B.1 and B.2, respectively.

In Appendix B.3 we solve the problem of a CARA investor and show that its solution coincides with the solution of our model only if one makes a linear approximation of a jump-related term in the HJB equation. Thus, the solution of our model is only an approximation to that of an investor with CARA preferences. However, we show that this approximation has a meaningful economic interpretation. Indeed, it corresponds to making the agent risk-neutral to regime-switching risk while remaining risk-averse to diffusion return-risk. In Appendix B.4, we formalize this argument by defining a set of 'source-dependent' recursive utility functions in which agents would have different aversion to jump risk and diffusion risk, building on Skiadas (2008), and Hugonnier, Pelgrin, and St-Amour (2012).

In Appendix B.5 we construct these preference specifications in a recursive framework, and show in Appendix B.6, that these preferences provide a micro-foundation to our objective functions for both the models in Section 2 and Section 4. Specifically our objective function is that of a recursive utility agent with source-dependent preferences that has constant absolute risk aversion towards return shocks, but has vanishing risk aversion to shocks driving the investment opportunity set (i.e., shocks to expected returns, variances, and transaction costs). We note that this 'source-dependent' utility function provides a micro-foundation to the objective function used both here and in Gârleanu and Pedersen (2013, GP). Effectively, the GP objective function is consistent with a preference-specification in which agents exhibit constant absolute risk aversion to the diffusion shocks driving stock returns, but are risk-neutral to shocks driving the return predicting factors. This can be proved straightforwardly following the same steps as in Appendix B.5 and Appendix B.6, but assuming that the expected return of stocks are driven by independent Brownian motions rather than by a regime switching model as in our framework. We leave the details to the interested reader.

B.1. Continuous-Time model of Price Changes: Infinite Horizon

We begin with a setting with N risky assets, in which the N-dimensional vector of price processes S_t, follows the process $dS_t = \mu(s_t)dt + \sigma(s_t)dZ_t$, driven by a K-dimensional vector of Brownian motions Z_t. The diffusion matrix $\sigma(s_t)$ is (N, K). $\mu(s_t)$ and $\Sigma(s_t) = \sigma(s_t)\sigma(s_t)^\top$ are, respectively, the N-vector of expected price changes and the $N \times N$ covariance matrix of price changes. Both μ and Σ are a function of a state variable s_t which follows a continuous-time Markov chain with transition intensities $\pi_{s,s'}$. For simplicity, we assume that the risk-free rate is zero, that there are only two states, and that the covariance matrix is full-rank.[23]

We consider the optimization problem of an agent with the following objective function:

$$\max_{\theta_t} \mathrm{E}\left[\int_{t=0}^{\infty} \left\{ n_t^\top \mu(s_t) - \frac{1}{2}\gamma n_t^\top \Sigma(s_t) n_t - \frac{1}{2}\theta_t^\top \Lambda(s_t)\theta_t \right\} e^{-\rho t} dt\right] \tag{A.1}$$

where n_t is the number of shares held by the investor and θ_t is the trading rate, that is to say, $dn_t = \theta_t dt$. We assume that Σ_s, Λ_s are real, symmetric, positive-definite matrices.[24]

We interpret this as an investor who maximizes his expected wealth, $\mathrm{E}[W_\tau]$ evaluated at a random time τ that is exponentially distributed with arrival intensity ρ. The wealth is accumulated capital gains net of quadratic trading and holding costs. Quadratic holding costs are standard in many papers (e.g, Duffie and Zhu (2017) and Du and Zhu (2017)). In these papers, holding costs are constant. In our framework, we assume that holding costs are proportional to the variance of the position held. For example, one can think of these holding costs as fees charged by a prime broker or as disutility perceived by a risk-averse investor for holding risky position. In the next section, we compare this set-up to that of a CARA investor facing quadratic transaction costs.

So with quadratic transaction and holding costs, the wealth dynamics are given by

$$dW_t = n_t dS_t - \frac{1}{2}\theta_t^\top \Lambda(s_t)\theta_t dt - \frac{1}{2}\gamma n_t^\top \Sigma(s_t) n_t dt, \tag{A.2}$$

$$dn_t = \theta_t dt. \tag{A.3}$$

Note that

$$\mathrm{E}[W_\tau] = \int_0^\infty \rho e^{-\rho t} W_t dt = W_0 + \int_0^\infty e^{-\rho t} dW_t$$

if the transversality condition $\lim_{T\to\infty} \mathrm{E}[e^{-\rho T} W_T] = 0$ holds. We solve this problem using dynamic programming. Define the value function:

[23] In our continuous time formulation, the price process is continuous despite the fact that there may be unpredictable jumps in means and covariances. While this type of price process can be supported in general equilibrium (e.g., in a Lucas-Breeden exchange economy with a representative log-investor and an aggregate output process that follows such a process), it would be interesting to consider an extension to the case where the stock prices can also experience jumps in their levels upon a regime shift. We leave such an extension for future research.

[24] Naturally we want $\theta^\top \Lambda \theta > 0 \; \forall \; \theta \neq 0$. Further, we have $\theta^\top \Lambda \theta = \frac{1}{2}\theta^\top \Lambda \theta + \frac{1}{2}(\theta^\top \Lambda \theta)^\top = \theta^\top (\frac{1}{2}\Lambda + \frac{1}{2}\Lambda^\top)\theta$. So if Λ is not symmetric we can replace it with $\frac{1}{2}(\Lambda + \Lambda^\top)$ which is.

$$J^s(n_t) = \max_{\theta_u} \mathrm{E}\left[\int_{u=t}^{\infty} e^{-\rho(u-t)}\left\{n_u^\top \mu(s_u) - \frac{1}{2}\gamma n_u^\top \Sigma(s_u) n_u - \frac{1}{2}\theta_u^\top \Lambda(s_u)\theta_u\right\} du \mid s_t = s\right]. \quad \text{(A.4)}$$

The HJB equation is

$$0 = \max_{\theta}\left\{n^\top \mu_s - \frac{1}{2}\gamma n^\top \Sigma_s n - \frac{1}{2}\theta^\top \Lambda_s \theta + (\nabla J^s)^\top \theta + \pi_{s,s'}(J^{s'} - J^s) - \rho J^s\right\},$$

where ∇J^s is the gradient of $J^s(n)$ and we use the subscript notation to denote the 'realization' of the matrix valued process in a particular state, i.e., $M(s_t)|_{s_t=s} = M_s$. The first order condition is given by

$$\theta = \Lambda_s^{-1} \nabla J^s.$$

Plugging this back into the HJB equation we get:

$$0 = n^\top \mu_s - \frac{1}{2}\gamma n^\top \Sigma_s n + \frac{1}{2}(\nabla J^s)^\top \Lambda_s^{-1} \nabla J^s + \pi_{s,s'}(J^{s'} - J^s) - \rho J^s \quad \text{(A.5)}$$

We guess that the value function is of the form:

$$J^s(n) = -\frac{1}{2} n^\top Q_s n + n^\top q_s + c_s$$

for some symmetric positive-definite matrix Q_s, a vector q_s and a scalar c_s.

Plugging this guess into the HJB equation gives

$$\begin{aligned} 0 = &-\frac{1}{2} n^\top \{Q_s \Lambda_s^{-1} Q_s + \gamma \Sigma_s - \rho Q_s + \pi_{s,s'}(Q_{s'} - Q_s)\} n \\ &+ n^\top \{\mu_s - Q_s \Lambda_s^{-1} q_s - \rho q_s + \pi_{s,s'}(q_{s'} - q_s)\} + \frac{1}{2} q_s^\top \Lambda_s^{-1} q_s - \rho c_s + \pi_{s,s'}(c_{s'} - c_s), \end{aligned}$$

which is verified if Q_s, q_s, and c_s solve the following set of matrix equations:

$$\begin{aligned} 0 &= Q_s \Lambda_s^{-1} Q_s + \gamma \Sigma_s - \rho Q_s + \pi_{s,s'}(Q_{s'} - Q_s) \\ 0 &= \mu_s - Q_s \Lambda_s^{-1} q_s - \rho q_s + \pi_{s,s'}(q_{s'} - q_s) \\ 0 &= \frac{1}{2} q_s^\top \Lambda_s^{-1} q_s - \rho c_s + \pi_{s,s'}(c_{s'} - c_s). \end{aligned}$$

To interpret the optimal trading strategy, note that the value function is maximized at the optimal aim portfolio $aim_s = Q_s^{-1} q_s$. Since $\nabla J^s = -Q_s n + q_s$ the optimal trading rate can be written as:

$$\theta = \Lambda_s^{-1} \nabla J^s = \Lambda_s^{-1} Q_s (aim_s - n)$$

So with the definition of trade intensity $\tau_s = \Lambda_s^{-1} Q_s$ we get the optimal position:

$$dn_t = \tau_s (aim_s - n_t) dt$$

with the same interpretation as in the discrete case.

B.2. Continuous time model of price changes: Finite horizon

To simplify the comparison to the expected utility framework provided in Appendix B.3, we consider now the finite horizon optimization problem. In this case, the agent has the following objective function:

$$\max_{\theta_t} \mathrm{E}\left[\int_{t=0}^{T} \left\{ n_t^\top \mu(s_t) - \frac{1}{2}\gamma n_t^\top \Sigma(s_t) n_t - \frac{1}{2}\theta_t^\top \Lambda(s_t)\theta_t \right\} dt \right] \tag{A.6}$$

where n_t are the number of shares held by the investor and θ_t is the trading rate, that is $dn_t = \theta_t dt$.

We interpret this as an investor who maximizes his expected wealth, $\mathrm{E}[W_T]$ evaluated at a fixed time T and faces quadratic transaction costs and holding costs, so that her wealth dynamics are given by:

$$dW_t = n_t dS_t - \frac{1}{2}\theta_t^\top \Lambda(s_t)\theta_t dt - \frac{1}{2}\gamma n_t^\top \Sigma(s_t) n_t dt \tag{A.7}$$

$$dn_t = \theta_t dt \tag{A.8}$$

Define the value function:

$$J^s(n_t, t) = \max_{\theta_u} \mathrm{E}\left[\int_{u=t}^{T} \left\{ n_u^\top \mu(s_u) - \frac{1}{2}\gamma n_u^\top \Sigma(s_u) n_u - \frac{1}{2}\theta_u^\top \Lambda(s_u)\theta_u \right\} du \,|\, s_t = s \right] \tag{A.9}$$

The HJB equation is

$$0 = \max_{\theta} \left\{ n^\top \mu_s - \frac{1}{2}\gamma n^\top \Sigma_s n - \frac{1}{2}\theta^\top \Lambda_s \theta + (\nabla J^s)^\top \theta + \pi_{s,s'}(J^{s'} - J^s) + \dot{J}^s \right\}$$

where ∇J^s is the gradient of $J^s(n,t)$ with respect to n and $\dot{J}^s$ is the partial derivative with respect to time.

The first order condition is

$$\theta = \Lambda_s^{-1} \nabla J^s$$

Plugging back into the HJB equation we get:

$$0 = n^\top \mu_s - \frac{1}{2}\gamma n^\top \Sigma_s n + \frac{1}{2}(\nabla J^s)^\top \Lambda_s^{-1} \nabla J^s + \pi_{s,s'}(J^{s'} - J^s) + \dot{J}^s \tag{A.10}$$

We guess that the value function is of the form:

$$J^s(n,t) = -\frac{1}{2} n^\top Q_s(T-t)n + n^\top q_s(T-t) + c_s(T-t)$$

for some symmetric positive-definite matrices Q_s, and a vector q_s and scalar c_s that are deterministic functions of time to maturity.

Plugging into the HJB equation we obtain

$$\begin{aligned} 0 = &-\frac{1}{2} n^\top \{Q_s \Lambda_s^{-1} Q_s + \gamma \Sigma_s - \dot{Q}_s + \pi_{s,s'}(Q_{s'} - Q_s)\} n \\ &+ n^\top \{\mu_s - Q_s \Lambda_s^{-1} q_s - \dot{q}_s + \pi_{s,s'}(q_{s'} - q_s)\} + \frac{1}{2} q_s^\top \Lambda_s^{-1} q_s - \dot{c}_s + \pi_{s,s'}(c_{s'} - c_s) \end{aligned} \tag{A.11}$$

We see that the guess is verified if Q_s, q_s, c_s solve the following set of ODE:

$$\dot{Q}_s = Q_s \Lambda_s^{-1} Q_s + \gamma \Sigma_s - Q_s + \pi_{s,s'}(Q_{s'} - Q_s) \tag{A.12}$$

$$\dot{q}_s = \mu_s - Q_s \Lambda_s^{-1} q_s - q_s + \pi_{s,s'}(q_{s'} - q_s) \tag{A.13}$$

$$\dot{c}_s = \frac{1}{2} q_s^\top \Lambda_s^{-1} q_s - c_s + \pi_{s,s'}(c_{s'} - c_s) \tag{A.14}$$

subject to boundary conditions $Q_s(0) = 0$, $q_s(0) = 0$, and $c_s(0) = 0$.

To interpret the optimal trading strategy, note that the value function is maximized at the optimal aim portfolio $aim_s(t) = Q_s^{-1}(T-t)q_s(T-t)$. Since $\nabla J^s = -Q_s n + q_s$ the optimal trade can be written as:

$$\theta = \Lambda(s)^{-1} \nabla J^s = \Lambda_s^{-1} Q_s (aim_s - n)$$

So with the definition of trade intensity $\tau_s(t) = \Lambda_s^{-1} Q_s(T-t)$ we get the optimal position:

$$dn_t = \tau_s(t)(aim_s(t) - n_t)dt$$

with the same interpretation as in the discrete case.

B.3. *The expected utility CARA investor*

Instead of assuming a risk-neutral investor with quadratic holding costs for the risky position, here we assume a CARA investor with wealth dynamics given by:

$$dW_t = n_t dS_t - \frac{1}{2}\theta_t^\top \Lambda(s_t)\theta_t dt \tag{A.15}$$

$$dn_t = \theta_t dt \tag{A.16}$$

We consider the optimization problem of an agent with the following objective function:

$$\max_{\theta_t} \mathrm{E}\left[-e^{-\gamma W_T}\right] \tag{A.17}$$

Define the 'value' function:

$$J^s(n_t, t) = \max_{\theta_u} \mathrm{E}\left[-e^{-\gamma \int_t^T dW_u} \mid s_t = s\right] \tag{A.18}$$

We seek a process θ such that $e^{-\int_0^t \gamma dW_u} J^s(n_t)$ is a supermartingale and a martingale at the optimal θ, that is:

$$0 = \max_\theta \left\{ -\gamma J^s(n^\top \mu_s - \frac{1}{2}\gamma n^\top \Sigma_s n - \frac{1}{2}\theta^\top \Lambda_s \theta) + (\nabla J^s)^\top \theta + \pi_{s,s'}(J^{s'} - J^s) + \dot{J}^s \right\}$$

where ∇J^s is the gradient w.r.t n of $J^s(n,t)$ and $\dot{J}^s$ denotes the time derivative.

The first order condition is

$$\theta = -\Lambda_s^{-1}\frac{\nabla J^s}{\gamma J^s}$$

Plugging back into the 'HJB equation' we get:

$$0 = -\gamma J^s(n^\top \mu_s - \frac{1}{2}\gamma n^\top \Sigma_s n) - \frac{1}{2}(\nabla J^s)^\top \Lambda_s^{-1}\frac{\nabla J^s}{\gamma J^s} + \pi_{s,s'}(J^{s'} - J^s) + \dot{J}^s$$

To solve we first do the transformation $J^s(n,t) = -e^{-\gamma j^s(n,t)}$ to obtain:

$$0 = n^\top \mu_s - \frac{1}{2}\gamma n^\top \Sigma_s n + \frac{1}{2}(\nabla j^s)^\top \Lambda_s^{-1} \nabla j^s + \pi_{s,s'} \frac{(1 - e^{-\gamma(j^{s'} - j^s)})}{\gamma} + \dot{j}^s$$

This system of equation needs to be solved numerically. But, by using the approximation $\frac{(1-e^{-\gamma x})}{\gamma} \approx x$ (valid for small γ) the equation simplifies to:

$$0 = n^\top \mu_s - \frac{1}{2}\gamma n^\top \Sigma_s n + \frac{1}{2}(\nabla j^s)^\top \Lambda_s^{-1} \nabla j^s + \pi_{s,s'}(j^{s'} - j^s) + \dot{j}^s$$

Comparing with the HJB equation in (A.10) obtained in the previous section, we see that the two equations are identical. Thus, the solution to the approximated HJB equation is identical to that obtained in the previous section, namely:

$$j^s(n,t) = -\frac{1}{2} n^\top Q_s(T-t) n + n^\top q_s(T-t) + c_s(T-t)$$

for some symmetric positive-definite matrices Q_s, and a vector q_s and scalar c_s all deterministic functions of time to maturity that solve the same system of ODEs given in equations (A.12)-(A.14) previously.

We note that as in the previous section, the value function is maximized at the aim portfolio given by $aim_s(t) = Q_s(T-t)^{-1} q_s(T-t)$ and that, since $-\Lambda_s^{-1} \frac{\nabla J^s}{\gamma J^s} = \Lambda_s^{-1} \nabla j^s = \Lambda_s^{-1} Q_s(T-t)(aim_s(t) - n)$ we can rewrite the optimal trading rule, with $\tau_s(t) = \Lambda_s^{-1} Q_s(T-t)$ as:

$$dn_t = \tau_s(t)(aim_s(t) - n_t)dt$$

So we see a 'rationalization' for our previous model. Essentially, the approximation in the CARA framework, boils down to making the agent risk-neutral with respect to regime shift shocks, while making her risk-averse with respect to the Brownian shocks. It turns out we can formalize this and derive a set of recursive preferences that are consistent with this behavior. These preferences belong to the source-dependent risk-aversion preferences developed in Hugonnier, Pelgrin, and St-Amour (2012), as we formalize next. The difference relative to their setting is that they work with recursive utility of a consumption flow and relative risk-aversion, whereas we develop the model for consumption at one final date and arbitrary utility functions (that include the absolute risk-aversion and relative risk-aversion case).

B.4. 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 Brownian motion Z_t and a Poisson process N_t with intensity λ_t, and who has expected utility of terminal wealth with twice-differential, increasing and concave utility function $U(W_T)$. Note that by definition $M_t = \mathrm{E}_t[U(W_T)]$ is a martingale and therefore we may write:

$$dM_t = \sigma_M dZ_t + \eta_M(dN_t - \lambda_t dt)$$

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

$$dH_t = \mu_H dt + \sigma_H dZ_t + \eta_H(dN_t - \lambda_t dt) \tag{A.19}$$

Then we have

$$dU(H_t) = [\frac{1}{2}U''(H)\sigma_H^2 + U'(H)(\mu_H - \lambda\eta_H)]dt + U'(H)\sigma_H dZ_t + (U(H+\eta_H) - U(H))dN_t$$

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

$$\mu_H = -\frac{1}{2}\frac{U''(H)}{U'(H)}\sigma_H^2 + \lambda(\eta_H - \frac{U(H+\eta_H) - U(H)}{U'(H)}) \tag{A.20}$$

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, η_H) of a backward-stochastic differential equation:

$$H_t = \mathrm{E}_t[W_T - \int_t^T \mu_H(H_t, \sigma_H, \eta_H)dt] \tag{A.21}$$

where the driver of the BSDE is given in equation (A.20) above.

To summarize, we have shown that, for an agent with an arbitrary wealth process W_t (driven by Brownian and Poisson shocks) who has expected utility of terminal wealth $\mathrm{E}[U(W_T)]$, we can define his certainty equivalent H_t in two different ways. First, the traditional definition $H_t = U^{-1}(\mathrm{E}_t[U(W_T)])$. Second, as the solution of the BSDE given in (A.20-A.21) above. Both are perfectly equivalent definitions. 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 sources of risk (in our case, to diffusion versus jump risk).

Specifically, we define the certainty equivalent of our "source-dependent stochastic differential utility" agent who consumes only at maturity T, as the solution (H_t, σ_H, η_H) of the following BSDE:

$$H_t = \mathrm{E}_t[W_T - \int_t^T \mu_H(H_t, \sigma_H, \eta_H)dt] \tag{A.22}$$

$$\mu_H = -\frac{1}{2}\frac{U_1''(H)}{U_1'(H)}\sigma_H^2 + \lambda\left(\eta_H - \frac{U_2(H+\eta_H) - U_2(H)}{U_2'(H)}\right) \tag{A.23}$$

where two different (twice-differential, strictly increasing and concave) utility functions U_i $i = 1, 2$ apply to the different sources of (e.g., diffusion versus jump) risk. In the next section we provide a heuristic derivation of this recursive utility based on a specific source-dependent certainty equivalent. This is similar to Skiadas (2008) and Hugonnier, Pelgrin, and St-Amour (2012).[25]

B.5. 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. Since wealth is driven by Z_t and N_t, we assume that prior to t, the certainty equivalent has dynamics as in (A.19). At any time $t < T$ the certainty equivalent is defined by the following recursion

$$W(H_t, 0, 0) = \mathrm{E}_t[W(H_t + \mu_H dt, \sigma_H dZ_t, \eta_H(dN_t - \lambda)dt)] \tag{A.24}$$

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

$$W(z_0, z_1, z_2) = U_1(z_0 + z_1) + \frac{U_1'(z_0)}{U_2'(z_0)}\left(U_2(z_0 + z_2) - U_2(z_0)\right)$$

[25] The only difference is that we do not have any intermediate consumption, and do not restrict to CRRA utility functions.

Using this we can rewrite the recursion (A.24), using the Itô rule for the right-hand side as:

$$U_1(H_t) = U_1(H_t) + U_1'(H_t)\mu_H dt + \frac{1}{2}U_1''(H_t)\sigma_H^2 dt - U_1'(H_t)\left(\eta_H - \frac{U_2(H_t+\eta_H) - U_2(H_t)}{U_2'(H_t)}\right)\lambda dt$$

Simplifying and rewriting we obtain the driver μ_H of the BSDE which defines the source-dependent SDU in equation (A.22) above.

B.6. *Source-Dependent SDU with Vanishing Jump Risk Aversion*

We will now show that our objective function in Appendix B.2 (that is the continuous-time version of our discrete-time model in Section 2) arises from source-dependent SDU with $U_i(x) = -e^{-\gamma_i x}$ for $i = 1, 2$ and letting $\gamma_2 \to 0$. Indeed, from the definition in (A.22) the certainty equivalent of such an agent is the solution (H, σ_H, η_H) of the BSDE:

$$H_t = \mathrm{E}_t\left[W_T - \int_t^T \left\{\frac{1}{2}\gamma_1\sigma_H^2 + \lambda(\eta_H - \frac{1-e^{-\gamma_2\eta_H}}{\gamma_2})\right\} du\right]$$

Now, if we let $\gamma_2 \to 0$ then the certainty equivalent is the solution (H, σ_H) of the BSDE:

$$H_t = \mathrm{E}_t\left[W_T - \int_t^T \frac{1}{2}\gamma_1\sigma_H^2 du\right]$$

Further, with W_t dynamics given in (A.15) above, we guess that the solution is of the form $H_t = W_t + J^s(n_t, t)$. Plugging this guess into the BSDE we find it is indeed a solution if $J^s(n_t, t)$ satisfies (note that this guess also implies that the diffusion of H is equal to the diffusion of W, that is $\sigma_H = \sigma_W$):

$$J^s(n_t, t) = \mathrm{E}_t\left[\int_t^T (n_u\mu(s_u) - \frac{1}{2}\gamma_1 n_u^\top\Sigma(s_u)n_u - \frac{1}{2}\theta_u^\top\Lambda(s_u)\theta_u)du\right]$$

This is the objective function we considered in Appendix B.2.

B.6.1. *Random horizon*

We can extend our analysis to the case of a random horizon τ with intensity ρ and define the certainty equivalent utility index for our source-dependent risk-averse investor (who has vanishing aversion to jump risk) as the solution (H_t, σ_H) of the BSDE:

$$\begin{aligned} H_t &= E_t[W_\tau - \int_t^\tau \frac{1}{2}\gamma\sigma_H^2 du] \\ &= E_t[W_t + \int_t^\tau \{dW_u - \frac{1}{2}\gamma\sigma_H^2\} du] \\ &= W_t + E_t[\int_t^\infty e^{-\rho u}\{dW_u - \frac{1}{2}\gamma\sigma_H^2\} du] \end{aligned}$$

subject to the appropriate transversality condition. The solution to this expression is a continuous time version of the models we consider in Section 2 and Section 4 of the paper depending on the assumption we make on the stock price process. Specifically if we assume wealth dynamics of the form $dW_t = n_t dS_t - \frac{1}{2}\theta^\top \Lambda \theta$ we obtain a continuous-time version of our model of price changes in Section 2, and if we consider $dW_t = x_t \frac{dS_t}{S_t} - \frac{1}{2}u^\top \Lambda u$ we obtain a continuous-time version of the model of dollar returns considered in Section 4.