NBER WORKING PAPER SERIES

LIQUIDITY REGIMES AND OPTIMAL DYNAMIC ASSET ALLOCATION

Pierre Collin-Dufresne
Kent D. Daniel
Mehmet Sa lam

Working Paper 24222
<http://www.nber.org/papers/w24222>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2018, Revised October 2018

For valuable comments and suggestions, we thank Darrell Duffie, Hui Guo, Ron Kaniel, an anonymous referee, and seminar participants at Princeton University, UCLA, University of Cincinnati and the 2018 Zurich Workshop on Asset Pricing. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w24222.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Pierre Collin-Dufresne, Kent D. Daniel, and Mehmet Sa lam. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Liquidity Regimes and Optimal Dynamic Asset Allocation
Pierre Collin-Dufresne, Kent D. Daniel, and Mehmet Sa lam
NBER Working Paper No. 24222
January 2018, Revised October 2018
JEL No. D53,G11,G12

ABSTRACT

We solve a portfolio choice problem when expected returns, volatilities and trading-costs follow a regime-switching model. The optimal policy trades towards an aim portfolio given by a weighted-average of the conditional mean-variance portfolios in all future states. The trading speed is higher in more persistent, riskier and higher-liquidity states. It can be optimal to overweight low Sharpe-ratio assets such as Treasury bonds because they remain liquid even in crisis states. We illustrate our methodology by constructing an optimal US equity market timing portfolio based on an estimated regime-switching model and on trading costs estimated using a large-order institutional trading dataset.

Pierre Collin-Dufresne
Ecole Polytechnique Federale de Lausanne
CDM SFI SFI-PCD
EXTRA 209
CH-1015 Lausanne
Switzerland
pierre.collin-dufresne@epfl.ch

Mehmet Salam
Assistant Professor of Finance
Carl H. Lindner College of Business
University of Cincinnati
408 Lindner Hall
Cincinnati, OH 45221
mehmet.saglam@uc.edu

Kent D. Daniel
Graduate School of Business
Columbia University
3022 Broadway, Uris Hall 421
New York, NY 10027
and NBER
kd2371@columbia.edu

1 Introduction

Mean-variance efficient portfolio optimization, introduced by Markowitz (1952), is still widely used in practice and taught in business schools. When either expected returns or the covariance matrix of returns changes over time then so will the conditional mean-variance efficient ‘Markowitz’ portfolio. In the presence of transaction costs however, it will generally not be optimal for investors to constantly rebalance to perfectly track the Markowitz portfolio. In recognition of this fact, practitioners generally employ ad-hoc adjustments to Markowitz optimization, but it is recognized that these approaches are not optimal (Grinold and Kahn 1999).

In a recent paper, Gârleanu and Pedersen (2013, GP) show that in the presence of quadratic transaction costs,¹ an investor with mean-variance preferences should adopt a trading rule that only partially rebalances from her current position towards an *aim portfolio* at a fixed *trading speed*.² They derive closed-form expressions for both the optimal aim portfolio and the trading speed that depend on the dynamics of expected returns, the quantity of and aversion to risk, and the magnitude of price impact. However, the GP model assumes that both the covariance matrix of price changes and the price-impact parameters are constant. In this paper we derive a closed-form solution for the optimal portfolio trading rule in a similar setting but where, in addition to expected returns, volatilities and transaction costs may be stochastic. This is consistent with considerable empirical evidence that stock return volatilities are stochastic and that transaction costs covary with the level of stock volatility (going back at least to Rosenberg (1972) for the former and to Stoll (1978) for the latter).

Here, we develop a closed-form solution for the optimal dynamic portfolio when expected returns, covariances, and price impact parameters follow a multi-state Markov switching model.³ Consistent with GP, we assume that the investor’s objective function is ‘dynamic’ mean-variance: investors maximize the expected discounted sum of portfolio returns net of trading costs, minus a penalty for the variance of portfolio returns. While this objective function is frequently used in practice and in academic papers, to our knowledge it lacks decision-theoretic foundations. In Appendices D-G, we show that the resulting optimal strategy corresponds to that of an agent who is risk-neutral to the regime-switching risk associated with variation in the investment opportunity set, while remaining risk-averse to the diffusion risk associated with return shocks. In Appendix E, we formalize this argument by defining a set of ‘source-dependent’ recursive (stochastic differential) utility preferences which result in the same objective function, building on Skiadas (2008), and Hugonnier, Pelgrin, and St-Amour (2012).⁴

In this setting, and for an agent with these preferences, we show that the optimal trading

¹Quadratic transaction costs emerge with a linear price impact model, i.e., trading Δ shares of a stock move its average price by $\lambda\Delta$ for a given constant λ .

²Litterman (2005) makes a similar point in an unpublished note.

³One additional advantage of this approach, which we discuss more below, and develop fully in Section 4, is that it allows us to specify a process for *return* volatility as opposed to *price-change* volatility in each state, consistent with the observed log-normality of prices.

⁴We thank the anonymous referee for encouraging us to provide decision-theoretic foundations to this objective function.

rule is similar to that derived in GP, namely to partially trade from the current position towards an aim portfolio. In GP the aim portfolio and trading speed are static. Here, when risk and trading costs can change, both the aim portfolio and the trading speed are conditional on the state. Specifically, the aim portfolio is a weighted average of the state-contingent Markowitz portfolios in all possible future states, where the weight on each conditional-Markowitz portfolio is a function of the likelihood of transitioning to that state, the state persistence, and the risk and transaction costs faced in that state relative to the current one. Similarly, the optimal trading speed depends on the relative magnitude of the transaction costs in various states and their transition probabilities. Moreover, while we solve the model in a discrete-time setting in the body of the paper, Appendices B and C solve a continuous time version of the model, and obtain consistent solutions.

To develop some of the intuition underlying the model, consider a simple setting with a single risky asset, and with two states: a low volatility state L where transaction costs are zero, and a high volatility state H where transaction costs are positive. When the economy is in the L -state, it is clearly optimal to trade (at infinite speed) all the way to the aim portfolio because transaction costs in that state are zero. In contrast, trading speed will be finite in the H -state. Further, the aim portfolio in the H -state will equal the conditional Markowitz portfolio in that state.⁵ Intuitively, in the H -state the investor should put zero-weight on the L -state Markowitz portfolio, because when the economy enters that state she will face no cost to rebalance (to the optimal aim portfolio) in that state. However, the aim portfolio in the L -state will be a weighed average of both H - and L -conditional Markowitz portfolios, and where the weight on the H -conditional Markowitz portfolio increases with the likelihood of transitioning from L to H , the persistence of the state H , and with the ratio of the volatilities in the H - and L -states.

One immediate implication of our model is that the aim portfolio will deviate significantly from the Markowitz benchmark in anticipation of possible future shifts in relative risk and/or transaction costs. Consider two assets, which can be thought of as a “Treasury” and a “Corporate” bond portfolio. Suppose that in the low-volatility state (state L), the Corporate portfolio has a far higher Sharpe ratio than the Treasury portfolio, so that the conditional Markowitz portfolio has most of its weight on Corporates. However, if the economy transitions to state H , then risk and trading costs will dramatically increase for Corporates, but will remain unchanged for Treasuries. In anticipation of this, the aim portfolio in the L -state will have a large Treasury position. Intuitively, if the economy transitions from the L to the H state, then the volatility of the Corporate portfolio will increase, its Sharpe ratio will fall, and it will become illiquid and costly to trade out of. Thus, the optimal dynamic “aim” portfolio preemptively reduces the position of the Corporate portfolio in the L -state.

A related implication of the model relates to the trading speed in the L state. When the current portfolio deviates from the aim portfolio, it may become optimal to trade the less liquid Corporate portfolio more aggressively than the Treasury portfolio. Intuitively, if the economy does transition to the H state, the Corporate portfolio will become much more expensive to trade, while

⁵That is, the aim portfolio in the H -state puts zero weight on the L -state Markowitz portfolio.

the Treasury portfolio will remain relatively liquid.

Our model also has implications for the popular (among practitioners) “risk-parity” strategy, which weights each asset class in such a way that each contributes an equal amount of volatility to the overall fund (see, e.g., Bridgewater (2011) and Asness, Frazzini, and Pedersen (2012)). Risk-parity can be thought of as the mean-variance efficient portfolio, when all asset classes have equal expected return and the correlations across asset classes are zero.⁶ Interestingly, even if it were optimal to hold a risk-parity portfolio at all times in the absence of transaction costs, we show that, when transaction costs and volatilities of various asset classes move over time in a correlated fashion, then it is optimal to deviate significantly from the risk-parity portfolio and that this deviation is larger in the low-risk regime. This is because the optimal portfolio in the low-risk regime, where transaction costs tend to be lowest, needs to put some weight on the optimal risk-parity portfolio in the high-risk regime, where high transaction costs will make it much more costly to delever out of the higher risk asset classes.

We present an empirical application of our framework in which a fund moves in and out of a stock market index, taking into account time varying expected returns, volatility and transaction costs. While our analytical results are all derived in the context of a regime-switching model of price changes (e.g., a “normal” model for prices), we show that our model remains tractable for a regime switching model of dollar returns (i.e., a “log-normal” model of prices). Since the latter model fits the data empirically better, we use this framework for the empirical implementation. We estimate a four state Markov regime switching model of returns and find, both in-sample and out of sample, evidence of time variation in first and second moments. To estimate the transaction cost parameters, we use a proprietary data set on realized trading costs incurred by a large financial institution trading on behalf of clients, as measured by the implementation shortfall of their trades (Perold 1988). We show that trading costs vary significantly across regimes, identified using the (highest) smoothed probabilities of the regimes. Not surprisingly, trading costs are higher for higher volatility regimes.

We test our trading strategy both in-sample and out-of-sample. For the out-of-sample test, the regime shifting model and the state probabilities are estimated using only data in the information set of an agent on the day preceding the trading date. We compare the performance of our optimal dynamic strategy to three alternatives: a constant dollar investment in the risky asset, corresponding to an unconditional estimate of the sample mean and variance of returns, a buy-and-hold policy that never trades and a myopic one-period mean-variance problem optimized for current transaction costs, but that ignores the future dynamics of the Markov regime switching model.

Our results show that the dynamic trading strategy significantly outperforms the other three strategies in the presence of transaction costs. To determine the source of the outperformance, we examine what source of time-variation leads to the biggest gains for the dynamic strategy. Specifically, we compare the gains obtained from timing changes in expected returns, in volatility, and in transaction costs. In this out-of-sample experiment, we find that the biggest benefits arise

⁶These assumptions are sometimes justified based on the difficulty to reliably estimate means and correlations.

from taking into account for time variation in market volatility and transaction costs, while the benefits from timing (estimated) variation in mean returns is more mixed. This reflects the fact that mean returns move less than one-for-one with variances. Our findings here are consistent with Moreira and Muir (2017), who show that there are gains to moving out of the market in response to an increase in market variance because the conditional market risk-premium moves less than one-for-one with its variance. Thus, since our model captures the time-variation in volatilities and the corresponding changes in transaction costs more accurately, it is able to manage the risk-exposure and the incurred transaction costs more reliably, which directly contributes to increasing the net performance.

There is large academic literature on portfolio choice that has extended Markowitz’s one period mean-variance setting to dynamic multiperiod setting with a time-varying investment opportunity set and more general objective functions.⁷ This literature has largely ignored realistic frictions such as trading costs, because introducing transaction costs and price impact in the standard dynamic portfolio choice problem tends to make it intractable. Indeed, most academic papers studying transaction costs focus on a very small number of assets (typically two), limited predictability, and typically no time-variation in second moments or transaction costs.⁸

Balduzzi and Lynch (1999) and Lynch and Balduzzi (2000) investigate the impact of fixed and proportional transaction costs on the utility costs and the optimal rebalancing rule of a single risky asset with time-varying expected return, using dynamic programming. 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 heuristic trading strategies and dual bounds for a general dynamic portfolio optimization problem with transaction costs and return predictability that can be applied to larger number of stocks. Longstaff (2001) studies a numerical solution to the one risky asset case with stochastic volatility when agents face liquidity constraints that force them to trade absolutely continuously.

Our paper is also related to the large literature on asset allocation under regime shifts. For example, Ang and Bekaert (2002) apply regime switching model to an international asset allocation problem to account for time-varying first and second moments of asset returns. Ang and Timmermann (2012) survey this literature in detail. One common observation in empirical papers estimating regimes is the low expected returns in high volatility states. Thus, these models would often suggest that the mean-variance investors should scale down their equity exposure in times of market stress. Our paper complements this literature by accounting for high transaction costs during these volatile periods. Jang, Keun Koo, Liu, and Loewenstein (2007) extend the models of Constantinides (1986) and Davis and Norman (1990) (e.g., one risky asset and one risk-free as-

⁷Merton (1969), Merton (1971), Brennan, Schwartz, and Lagnado (1997), Kim and Omberg (1996), Campbell and Viceira (2002), Campbell, Chan, and Viceira (2003), Liu (2007), Detemple and Rindisbacher (2010) and many more. See Cochrane (2007) for a survey.

⁸Constantinides (1986), Davis and Norman (1990), Dumas and Luciano (1991), Shreve and Soner (1994) study the two-asset (one risky-one risk-free) case with i.i.d. returns. Liu (2007) studies the multi-asset case under CARA preferences and *i.i.d.* returns. Cvitanić (2001) surveys this literature.

set) with regime-switching fundamental parameters. They consider a small investor with no price impact and illustrate that proportional transaction costs may have first-order effect on liquidity premia. In comparison, we consider a regime switching model in which an investor with price impact can trade multiple risky assets.

As noted earlier, our paper is most closely related to Litterman (2005) and Gârleanu and Pedersen (2013, GP). They obtain a closed-form solution for the optimal portfolio choice in a model where: (1) expected price change per share for each security is a linear, time-invariant function of a set of autoregressive predictor variables; (2) the covariance matrix of price changes is constant; (3) trading costs are a quadratic function of the number of shares traded, and (4) investors have a linear-quadratic objective function. Their approach relies heavily on linear-quadratic stochastic programming (see, *e.g.*, Ljungqvist and Sargent (2004)). Our approach uses a similar objective function, but allows for time-variation in means, volatilities, and transaction costs, albeit within a regime-switching framework. Moreover, in contrast with the GP framework, our framework is equally tractable when expected-price changes are constant in each state of the regime switching model (i.e., prices follow arithmetic Brownian motion) or when expected returns, conditional on the state, are constant (i.e., prices follow geometric Brownian motion). Since the latter is a more realistic description of historical returns, it is the one we use for our empirical implementation.

2 A Regime Switching Model for Price Changes

We begin with a setting with N risky assets, in which the N -dimensional vector of price changes from period t to $t + 1$, dS_t , follows the process:

$$\begin{aligned} \mathbb{E}[dS_t] &= \mu(s_t) \\ \mathbb{E}[(dS_t - \mu(s_t))(dS_t - \mu(s_t))^\top] &= \Sigma(s_t) \end{aligned}$$

$\mu(s_t)$ and $\Sigma(s_t)$ 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 Markov chain with transition probabilities $\pi_{s,s'}$. In Section 4, we will solve for the optimal dynamic strategy when returns, rather than price change, follow this process.

We consider the optimization problem of an agent with the following objective function:⁹

$$\max_{n_t} \mathbb{E} \left[\sum_{t=0}^{\infty} \rho^t \left\{ n_t^\top \mu(s_t) - \frac{1}{2} \gamma n_t^\top \Sigma(s_t) n_t - \frac{1}{2} \Delta n_t^\top \Lambda(s_t) \Delta n_t \right\} \right] \quad (1)$$

The agent chooses her holdings n_t in each period t so as to maximize this objective function.

⁹In the appendix, we provide two ways to micro-found this objective function. First, it corresponds to an agent who maximizes her expected terminal wealth $\mathbb{E}[W_\tau]$ at some random horizon τ , drawn from an exponential distribution with intensity $-\ln \rho$, and who faces quadratic holding costs that are proportional to the variance of the position as well as quadratic trading costs. Second, this corresponds to the certainty equivalent of an agent with source dependent stochastic differential utility who has CARA aversion with risk-aversion coefficient γ towards return shocks and vanishing risk-aversion $\gamma_2 \rightarrow 0$ towards regime-shifts.

Specifically, at the end of period $t - 1$, the agent hold n_{t-1} shares of the N assets. At this point the agent observes the state s_t , and trades $\Delta n_t = n_t - n_{t-1}$ shares. As noted earlier, consistent with GP we specify a linear price impact model. $\Lambda(s_t)$ is the price impact matrix, so the N -vector of price concessions is $\Lambda(s_t)\Delta n_t$ and the total (dollar) cost of trading in period t is therefore $\frac{1}{2}\Delta n_t^\top \Lambda(s_t)\Delta n_t$. We assume that Σ_s and Λ_s are real symmetric positive-definite matrices.¹⁰

This objective function in equation (1) is the same as that considered by GP, namely that of an investor who maximizes a discounted sum of mean-variance criterion in every period, net of trading costs. In the absence of transaction costs (when $\Lambda_s = 0$), the optimal solution would be to hold the conditionally mean-variance optimal Markowitz portfolio $m_s = (\gamma\Sigma_s)^{-1}\mu_s$ at all times. Further, if there were no time-variation in the investment opportunity set (that is if μ_s and Σ_s were constant), then it would be always optimal to hold the mean-variance efficient Markowitz portfolio. However, when there are transaction costs and the opportunity set is time-varying, it becomes optimal for the investor to rebalance the portfolio, and deviate from the conditionally mean-variance efficient portfolio.

In the GP framework, the conditional mean of stock price changes (μ_s) follows an AR(1) process, but the covariance matrix Σ and the matrix of transaction cost parameters Λ are required to be time invariant. In our framework Σ and Λ vary across states. Using a Markov regime switching model allows us to obtain tractable solutions even though the model is not in the standard linear quadratic framework.

For simplicity we begin by considering only a two-state Markov chain model, with states H and L , but we generalize this to more states in Section 2.4. We will use the following notation throughout: for all t where $s_t = s \in \{H, L\}$, $s_{t+1} = z \in \{H, L\}$ and $s' = \{H, L\} \setminus s$. Then, using the dynamic programming principle, the value function $V(n_{t-1}, s)$ satisfies

$$V(n_{t-1}, s) = \underset{n_t}{\text{maximize}} \left(n_t^\top \mu_s - \frac{1}{2} \Delta n_t^\top \Lambda_s \Delta n_t - \frac{\gamma}{2} n_t^\top \Sigma_s n_t + \rho \mathbb{E} [V_t(n_t, z)] \right).$$

We guess the following quadratic form for our value functions:

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

where Q_s is a symmetric $N \times N$ matrix and q_s, c_s are N dimensional vectors of constants for $s \in \{H, L\}$. We now define the expectation conditional on state s for any matrix M_s to be $\overline{M}_s = \pi_{s,s} M_s + \pi_{s,s'} M_{s'}$. With this notation, the right hand side of the HJB equation we are optimizing can be rewritten as a quadratic objective:

$$-\frac{1}{2} n_t^\top J_s n_t + n_t^\top j_s + k_s$$

¹⁰Naturally 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.

where

$$\begin{aligned} J_s &= \gamma \Sigma_s + \Lambda_s + \rho \bar{Q}_s \\ j_s &= \mu_s + \Lambda_s n_{t-1} + \rho \bar{q}_s \\ k_s &= -\frac{1}{2} n_{t-1} \Lambda_s n_{t-1} + \rho \bar{c}_s \end{aligned}$$

This is optimized for $n_t = J_s^{-1} j_s$, that is :

$$n_t = (\gamma \Sigma_s + \Lambda_s + \rho \bar{Q}_s)^{-1} (\mu_s + \rho \bar{q}_s + \Lambda_s n_{t-1})$$

Further, the optimized value is simply $\frac{1}{2} j_s^\top J_s^{-1} j_s + k_s$. Thus matching coefficients we find that the matrices Q_s, q_s for $s = H, L$ must satisfy the system of equations:

$$Q_s = -\Lambda_s (\gamma \Sigma_s + \Lambda_s + \rho \bar{Q}_s)^{-1} \Lambda_s + \Lambda_s, \quad (2)$$

$$q_s = \Lambda_s (\gamma \Sigma_s + \Lambda_s + \rho \bar{Q}_s)^{-1} (\mu_s + \rho \bar{q}_s). \quad (3)$$

Note that given a solution for Q_H and Q_L , we can obtain q_H and q_L in closed-form as a matrix weighted average of μ_H and μ_L . While we are not aware of a closed-form solution for Q_H and Q_L in general, it is straightforward to obtain a numerical solution to the coupled Riccati matrix equation, as we discuss in Lemma 2 below. Further, for a variety of special cases we consider below, it is possible to obtain closed-form solutions.

With a solution in hand, we can define the conditional aim portfolio as the portfolio that maximizes the value function at any time t conditional on the state. We can now characterize the optimal trading rule and the aim portfolios.

Theorem 1 *The optimal trade at time t in state s is a matrix weighted average of the current position vector and the conditional aim portfolio:*

$$n_t = (I - \tau_s) n_{t-1} + \tau_s \text{aim}_s \quad (4)$$

where the trading speed $\tau_s = I$ (and $Q_s = 0$) if $\Lambda_s = 0$, and else $\tau_s = \Lambda_s^{-1} Q_s \ \forall s = \{H, L\}$ where (Q_H, Q_L) solve a system of coupled equations:

$$I - \Lambda_s^{-1} Q_s = [\Lambda_s^{-1} (\gamma \Sigma_s + \rho \pi_{ss'} Q_{s'}) + I + \rho \pi_{ss} \Lambda_s^{-1} Q_s]^{-1} \quad (5)$$

The aim portfolio, which maximizes the value function conditional on the current state, is given by

$$\text{aim}_s = (\gamma \Sigma_s + \rho \bar{Q}_s)^{-1} (\mu_s + \rho \bar{q}_s) \quad (6)$$

Further, the aim portfolio is a weighted average of the conditional Markowitz portfolios ($m_s = (\gamma \Sigma_s)^{-1} \mu_s$):

$$\text{aim}_s = (I - \alpha_s) m_s + \alpha_s m_{s'} \ \forall s = H, L \quad (7)$$

where

$$\alpha_s = \{(\gamma + \rho\pi_{s's}Q_{s'}\Sigma_{s'}^{-1}Q_sQ_{s'}^{-1})\Sigma_s + \rho\pi_{ss'}Q_{s'}\}^{-1}\rho\pi_{ss'}Q_{s'}$$

Proof. Optimizing the value function with respect to n_t gives:

$$aim_s = (Q_s)^{-1}(q_s) \quad \forall s = H, L$$

Substituting from the definitions in equations (2) and (3) we obtain:

$$\begin{aligned} aim_s &= \left(-\Lambda_s(\gamma\Sigma_s + \Lambda_s + \rho\overline{Q}_s)^{-1}\Lambda_s + \Lambda_s\right)^{-1} \left(\Lambda_s(\gamma\Sigma_s + \Lambda_s + \rho\overline{Q}_s)^{-1}(\mu_s + \rho\overline{q}_s)\right) \\ &= \left(-(\gamma\Sigma_s + \Lambda_s + \rho\overline{Q}_s)^{-1}\Lambda_s + I\right)^{-1} (\gamma\Sigma_s + \Lambda_s + \rho\overline{Q}_s)^{-1}(\mu_s + \rho\overline{q}_s) \\ &= (\gamma\Sigma_s + \rho\overline{Q}_s)^{-1}(\mu_s + \rho\overline{q}_s) \end{aligned}$$

where the last equality obtains by noting that if we define the matrix

$$M = \left(-(\gamma\Sigma_s + \Lambda_s + \rho\overline{Q}_s)^{-1}\Lambda_s + I\right)^{-1} (\gamma\Sigma_s + \Lambda_s + \rho\overline{Q}_s)^{-1}$$

then

$$M^{-1} = (\gamma\Sigma_s + \Lambda_s + \rho\overline{Q}_s) \left(-(\gamma\Sigma_s + \Lambda_s + \rho\overline{Q}_s)^{-1}\Lambda_s + I\right) = (\gamma\Sigma_s + \rho\overline{Q}_s),$$

which immediately implies that $M = (\gamma\Sigma_s + \rho\overline{Q}_s)^{-1}$.

We then expand the expression for aim_s :

$$\begin{aligned} aim_s &= (\gamma\Sigma_s + \rho\pi_{ss}Q_s + \rho\pi_{ss'}Q_{s'})^{-1}(\mu_s + \rho\overline{Q}_s) \\ \Rightarrow (\gamma\Sigma_s + \rho\pi_{ss}Q_s + \rho\pi_{ss'}Q_{s'})aim_s &= (\gamma\Sigma_s m_s + \rho\pi_{ss}Q_s aim_s + \rho\pi_{ss'}Q_{s'} aim_{s'}) \\ \Rightarrow (\gamma\Sigma_s + \rho\pi_{ss'}Q_{s'})aim_s &= (\gamma\Sigma_s m_s + \rho\pi_{ss'}Q_{s'} aim_{s'}) \\ \Rightarrow aim_s &= (\gamma\Sigma_s + \rho\pi_{ss'}Q_{s'})^{-1}(\gamma\Sigma_s m_s + \rho\pi_{ss'}Q_{s'} aim_{s'}) \end{aligned}$$

We then substitute for $aim_{s'} = (\gamma\Sigma_{s'} + \rho\pi_{s's}Q_s)^{-1}(\gamma\Sigma_{s'} m_{s'} + \rho\pi_{s's}Q_s aim_s)$ and obtain after dividing by γ

$$\left[\Sigma_s + \frac{\rho}{\gamma}\pi_{ss'}Q_{s'} \left(I - (\gamma\Sigma_{s'} + \rho\pi_{s's}Q_s)^{-1}\rho\pi_{s's}Q_s\right)\right] aim_s = \Sigma_s m_s + \rho\pi_{ss'}Q_{s'} (\gamma\Sigma_{s'} + \rho\pi_{s's}Q_s)^{-1}\Sigma_{s'} m_{s'}.$$

Using the simple identity $I - (F + G)^{-1}G = (F + G)^{-1}F$, with $F = \gamma\Sigma_{s'}$ and $G = \rho\pi_{s's}Q_s$, we finally obtain

$$\{\Sigma_s + \rho\pi_{ss'}Q_{s'}[\gamma\Sigma_{s'} + \rho\pi_{s's}Q_s]^{-1}\Sigma_{s'}\} aim_s = \Sigma_s m_s + \rho\pi_{ss'}Q_{s'}[\gamma\Sigma_{s'} + \rho\pi_{s's}Q_s]^{-1}\Sigma_{s'} m_{s'}.$$

Thus, this shows that we can write $aim_s = (I - \alpha_s)m_s + \alpha_s m_{s'}$ where

$$\alpha_s = \{\Sigma_s + \rho\pi_{ss'}Q_{s'}[\gamma\Sigma_{s'} + \rho\pi_{s's}Q_s]^{-1}\Sigma_{s'}\}^{-1}\rho\pi_{ss'}Q_{s'}[\gamma\Sigma_{s'} + \rho\pi_{s's}Q_s]^{-1}\Sigma_{s'}$$

which can be further simplified to $\alpha_s = \{(\gamma + \rho\pi_{s's}Q_{s'}\Sigma_{s'}^{-1}Q_sQ_{s'}^{-1})\Sigma_s + \rho\pi_{ss'}Q_{s'}\}^{-1}\rho\pi_{ss'}Q_{s'}$.

■

Equation (4) shows that this optimal dynamic strategy is to trade to a portfolio with shares n_t that is a linear combination of the current portfolio n_{t-1} and of the aim portfolio aim_s . τ_s is the matrix that specifies how quickly the investor should trade towards the aim portfolio. $\tau_s = I$ means that, in state s , the investor should immediately and fully trade to aim_s . $\tau_s = 0$ means that the investor should not trade.

The state-contingent aim portfolio aim_s is defined as the portfolio that would maximize the value function in that state. Another interpretation of the aim portfolio is as the *no-trade* portfolio, i.e., the portfolio for which the optimal trade is zero, as long as the state does not change.¹¹ The speed at which we trade towards the aim portfolio is, in general, dependent on the state. That is, it is typically increasing in variance and decreasing in the transaction costs, which may be state dependent in our framework. In the case (similar to GP) where only expected returns are stochastic (and covariances and transaction costs are constant) the trading speed is constant as well. Further, the aim portfolio is state dependent. When either a state is absorbing ($\pi_{ss} = 1$) or transaction costs are zero ($\Lambda_s = 0$) then the aim portfolio is equal to the conditional mean-variance Markowitz portfolio (m_s). But in general, the aim portfolio is a weighted average of the conditional mean-variance portfolio across states, where the weight on each state is typically higher, if the variance of returns or the transaction cost is higher in that state.

We now consider a few special cases to gain further insights into the optimal trading rule.

2.1 The case where only μ_s changes with the state (GP)

If only μ_s changes with the state (i.e., if $\Sigma_s = \Sigma$ and $\Lambda_s = \Lambda$ for all s) then the solution $Q_s = Q$ is independent of the state and satisfies:

$$I - \Lambda^{-1}Q = [\gamma\Lambda^{-1}\Sigma + I + \rho\Lambda^{-1}Q]^{-1}$$

This equation has an explicit solution as we show in the following lemma.¹²

Lemma 1 *Consider the diagonalization of the matrix $\Lambda^{-1}\Sigma = F \text{diag}(\ell_i) F^{-1}$ in terms of its eigenvalues $\ell_i \forall i = 1, \dots, n$. Then note that*

$$I - F^{-1}\Lambda^{-1}QF = [\gamma \text{diag}(\ell_i) + I + \rho F^{-1}\Lambda^{-1}QF]^{-1}$$

¹¹Note that, because the vector of security holdings n has units of shares, and because the price change process is a function only of the state, the optimal portfolio will not change when prices change.

¹²We note that since Σ, Λ are assumed to be symmetric matrices with (strictly) positive real eigenvalues, then $\Lambda^{-1}\Sigma$ is diagonalizable. First, note that since Λ is real symmetric positive definite then so is its inverse. This implies we can decompose $\Lambda^{-1} = M^{\frac{1}{2}}M^{\frac{1}{2}}$. It follows that $M^{\frac{1}{2}}\Sigma M^{\frac{1}{2}}$ is symmetric and positive definite (as $x^{\top}M^{\frac{1}{2}}\Sigma M^{\frac{1}{2}}x = (M^{\frac{1}{2}}x)^{\top}\Sigma(M^{\frac{1}{2}}x) > 0 \forall x \neq 0$ since Σ is positive definite) and therefore has positive real eigenvalues. In turn, it is easy to show that $\Lambda^{-1}\Sigma = M^{\frac{1}{2}}M^{\frac{1}{2}}\Sigma$ has the same eigenvalues as $M^{\frac{1}{2}}\Sigma M^{\frac{1}{2}}$.

It follows that $Q = \Lambda F \text{diag}(\eta_i) F^{-1}$ such that the η_i solve the quadratic equations ($\forall i = 1, \dots, n$):

$$1 - \eta_i = [\gamma \ell_i + 1 + \rho \eta_i]^{-1}$$

that is:

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

This implies that the trading speed $\tau_s = \Lambda_s^{-1} Q_s = F \text{diag}(\eta_i) F^{-1}$ is independent of the state. That is, investors trade at a constant speed towards their aim portfolio independent of the state. The speed of trading for specific stock i is increasing in the agent's time discount rate and in the agent's risk-aversion. Furthermore, for the special case where Λ and Σ are diagonal matrices, then the speed of trading stock i is increasing in $\ell_i = \Sigma_{ii}/\Lambda_{ii}$, that is the ratio of a stock's variance to its cost of trading.

While the trading speed is constant, the aim portfolios differ across states. Indeed, using Theorem 1, the aim portfolio in state s can be computed as:

$$aim_s = (I - \alpha_s) m_s + \alpha_s m_{s'}$$

where

$$\begin{aligned} \alpha_s &= \{\gamma \Sigma + \rho \pi_{s's} Q + \rho \pi_{ss'} Q\}^{-1} \rho \pi_{ss'} Q \\ &= (\gamma Q^{-1} \Sigma + (\rho \pi_{s's} + \rho \pi_{ss'}) I)^{-1} \rho \pi_{ss'} \\ &= F \text{diag} \left(\frac{\rho \pi_{ss'}}{\gamma \ell_i / \eta_i + \rho \pi_{s's} + \rho \pi_{ss'}} \right) F^{-1} \end{aligned}$$

The state s aim portfolio is a weighted average of the conditional Markowitz portfolios in the current state (s) and in the alternative state (s'), where the weight on the current state Markowitz portfolio is increasing in the persistence of that state $\pi_{s,s}$ and in risk-aversion γ , but decreasing in the time discount factor ρ , and the persistence of the other state $\pi_{s',s'}$. Furthermore, the weight is also stock-specific and increasing for stock i in ℓ_i , which captures the notion that the more risky a stock is relative to its trading cost the more weight we should put on the conditional Markowitz portfolio for computing the aim portfolio.

To a large extent these results are consistent with the findings of GP, albeit with a different model of the time-variation in expected returns. The more interesting case is when we also allow covariances and transaction costs to change across states. In that case, both trading speed and aim portfolios change across states.

2.2 The case where $\Lambda_L = 0$ and $\Lambda_H > 0$

When transaction costs are zero in state L , then the solution implies $Q_L = 0$, and that Q_H solves a one-dimensional equation:

$$I - \Lambda_H^{-1}Q_H = [\gamma\Lambda_H^{-1}\Sigma_H + I + \rho\pi_{HL}\Lambda^{-1}Q_H]^{-1}$$

We note that this equation is identical to that obtained in the previous section with an adjusted time discount rate ($\rho\pi_{HL}$). It follows that the solution is

$$Q_H = \Lambda_H F_H \text{diag}(\eta_{H,i}) F_H^{-1},$$

where $(\ell_{H,i}, F_H)$ diagonalize the matrix $\Lambda_H^{-1}\Sigma_H = F_H \text{diag}(\ell_{H,i}) F_H^{-1}$ and the $\eta_{H,i}$ are given by:

$$\eta_{H,i} = \frac{\rho\pi_{HL} - 1 - \ell_{H,i}\gamma + \sqrt{(\rho\pi_{HL} - 1 - \ell_{H,i}\gamma)^2 + 4\ell_{H,i}\gamma\rho\pi_{HL}}}{2\rho\pi_{HL}}.$$

We can calculate the optimal trading speeds and the aim portfolios in both states. As discussed earlier, in the L state where transaction costs are zero, it is optimal to move instantaneously to the aim portfolio, that is $\tau_L = I$. In contrast, in the high transaction cost state H , it is optimal to trade slowly, with a trading speed $\tau_H = F_H \text{diag}(\eta_{H,i}) F_H^{-1}$, towards the aim portfolio. The aim portfolio in the high transaction cost state H is the conditional Markowitz portfolio, that is $aim_H = m_H = (\gamma\Sigma_H)^{-1}\mu_H$. Intuitively, in the state H , the aim portfolio does not take into account the investment opportunity set in the zero-transaction cost state L , because when the economy transitions to state L the investor can immediately rebalance to the first best position at zero cost. However, in the zero transaction cost state, the aim portfolio is a linear combination of the two Markowitz portfolios m_H and m_L : $aim_L = (I - \alpha_L)m_L + \alpha_L m_H$, where the weight put on the H -state Markowitz portfolio is $\alpha_L = [\gamma\Sigma_L + \rho\pi_{LH}Q_H]^{-1}\rho\pi_{LH}Q_H$. To summarize, when there are no transaction costs in the low state the optimal trading strategy is:

$$\begin{aligned} n_{H,t} &= (I - \tau_H)n_{t-1} + \tau_H m_H \\ \tau_H &= F_H \text{diag}(\eta_{H,i}) F_H^{-1} \\ n_{L,t} &= aim_L = (I - \alpha_L)m_L + \alpha_L m_H \\ \alpha_L &= [\gamma\Sigma_L + \rho\pi_{LH}Q_H]^{-1}\rho\pi_{LH}Q_H \end{aligned}$$

2.3 The case with $\Lambda_L > 0$ and $\Lambda_H = \infty$

We now consider the polar case, where transaction costs are infinite in the H -state. Clearly, it is then optimal not to rebalance in the high state. Following the derivation of our model, with no rebalancing in the H -state, we see that the equation for Q_H simplifies to:

$$Q_H = \gamma\Sigma_H + \rho\bar{Q}_H$$

In turn, this implies that the equation for Q_L becomes:

$$I - \Lambda_L^{-1}Q_L = [\gamma\Lambda_L^{-1}(\Sigma_L + \frac{\rho\pi_{LH}}{1-\rho\pi_{HH}}\Sigma_H) + I + \rho_L\Lambda_L^{-1}Q_L]^{-1}$$

with $\rho_L = \rho(\pi_{LL} + \frac{\rho\pi_{LH}}{1-\rho\pi_{HH}})$. This equation admits an explicit solution as before, in terms of the diagonalization of the matrix $\Lambda_L^{-1}(\Sigma_L + \frac{\rho\pi_{LH}}{1-\rho\pi_{HH}}\Sigma_H) = F_L \text{diag}(\ell_{L,i})F_L^{-1}$.

It follows that the solution is $Q_L = \Lambda_L F_L \text{diag}(\eta_{L,i})F_L^{-1}$ where the $\eta_{L,i}$ are given by:

$$\eta_{L,i} = \frac{\rho_L - 1 - \ell_{L,i}\gamma + \sqrt{(\rho_L - 1 - \ell_{L,i}\gamma)^2 + 4\ell_{L,i}\gamma\rho_L}}{2\rho_L}$$

Then the optimal trading strategy is:

$$n_{H,t} = n_{t-1}$$

$$n_{L,t} = (I - \Lambda_L^{-1}Q_L)n_{t-1} + \Lambda_L^{-1}Q_L \text{aim}_L$$

$$\text{aim}_L = (1 - \alpha_L)m_L + \alpha_L m_H$$

$$\alpha_L = \{(1 - \rho\pi_{HH})\Sigma_H^{-1}\Sigma_L + \rho\pi_{LH}\}^{-1}\rho\pi_{LH}$$

To summarize, when transaction costs are infinite in state H it is clearly optimal to not rebalance in that state. Instead, in state L , both the speed of trading and the aim portfolio depend on the investment opportunity set in the H state. The aim portfolio puts more weight on the H -conditional Markowitz portfolio the higher the probability to transition to that state (π_{LH}), the more persistent the state is (π_{HH}), and the higher the variance of returns in that state relative to the L -state ($\Sigma_H^{-1}\Sigma_L$). The trading speed on the other hand increases in both Σ_H and Σ_L as well as the persistence of the low and high states.

2.4 The general case

For the general case, we need to solve the system of coupled matrix equations (5) for (Q_H, Q_L) :

$$I - \Lambda_s^{-1}Q_s = [\Lambda_s^{-1}(\gamma\Sigma_s + \rho\pi_{ss'}Q_{s'}) + I + \rho\pi_{ss}\Lambda_s^{-1}Q_s]^{-1}$$

While we cannot solve the system in general, we observe that in the special case where the eigenvectors of the covariance and transaction cost matrices remain identical across states and only the eigenvalues change, the system does admit a simple explicit solution. This is a ‘knife-edge case’ in the general space of unconstrained matrices.¹³ Still, it is an interesting parametrization, as it nests the special case where both the transaction cost and covariance matrices are diagonal with arbitrary coefficients in all states. It also nests the special case considered in GP where the transaction cost

¹³To understand the parameter restrictions, note that given that both transaction cost and covariance matrices are symmetric positive definite they each would have $n(n+1)/2$ free parameters (subject to the restriction that they are positive definite). When we constrain all 4 (i.e., 2 in each state) matrices to have the same eigenvectors, then the total number of free parameters become $n(n+1)/2$ parameters for one matrix and only n parameters for the other 3 matrices. Indeed, since the latter matrices inherit the eigenvectors of the first matrix, each has only n free parameters, corresponding to their positive eigenvalues.

matrix is proportional to the covariance matrix, but here with possible state-dependent constants of proportionality (i.e., where $\Lambda_s = \lambda_s \Sigma_s$ and $\Sigma_{s'} = \beta \Sigma_s$ for some positive scalars $\beta, \lambda_s, \lambda_{s'}$). Also, for the general case of unconstrained matrices that can be solved numerically, we propose a simple and efficient algorithm to compute the solution.

We summarize these results in the following

Lemma 2 *If $\Lambda_s = F \text{diag}(\lambda_{i,s}) F^{-1}$ and $\Sigma_s = F \text{diag}(v_{s,i}) F^{-1} \forall s = H, L$ then the solution of the system of matrix equations (5) is $Q_s = \Lambda_s F \text{diag}(\eta_{s,i}) F^{-1}$ where $\forall i = 1, \dots, n$ the constants $(\eta_{H,i}, \eta_{L,i})$ solve the system of coupled quadratic equations:*

$$\frac{\lambda_{i,s}}{1 - \eta_{i,s}} = \gamma v_{i,s} + \rho \pi_{ss'} \eta_{i,s'} \lambda_{i,s'} + \lambda_{i,s} + \rho \pi_{ss} \eta_{i,s} \lambda_{i,s}$$

In general, when Σ_s, Λ_s do not have identical eigenvectors across states, then the solution to the system of matrix equations (5) can be obtained by the following recursion.

Given an initial (Q_H^{n-1}, Q_L^{n-1}) , perform the eigenvalue decomposition (for $s = H, L$) of $\Lambda_s^{-1}(\gamma \Sigma_s + \rho \pi_{ss'} Q_{s'}^{n-1}) = F_s \text{diag}(\ell_{i,s}) F_s^{-1}$. Then set $Q_s^n = \Lambda_s F_s \text{diag}(\eta_{i,s}) F_s^{-1}$ where the $\eta_{i,s}$ solve the equation

$$1 - \eta_{i,s} = [\ell_{i,s} + 1 + \rho \pi_{ss} \eta_{i,s}]^{-1},$$

that is:

$$\eta_{i,s} = \frac{\rho \pi_{ss} - 1 - \ell_{i,s} + \sqrt{(\rho \pi_{ss} - 1 - \ell_{i,s})^2 + 4 \ell_{i,s} \rho \pi_{ss}}}{2 \rho \pi_{ss}},$$

and iterate until convergence. It is natural to use as an initial guess for Q_s^0 either the zero matrix, or the solution corresponding to $\pi_{ss} = 1$.

We conjecture that the algorithm will be especially useful for large number of stocks, where iterating over the $N(N+1)$ elements of the Q_L and Q_H matrices should be less efficient than iterating over the $2N$ diagonal $\eta_{i,s}$ elements. In our applications, we found that only three to five iterations are sufficient to achieve convergence. Given a numerical solution of the Q_H and Q_L matrices, we can analyze the optimal trading rule and aim portfolios.

3 Implications of the Model

In this section, we illustrate the insights of our model using two simple numerical experiments. In the first application, we have two assets differing in their ranking of Sharpe ratios across two states of the economy. We analyze the aim portfolio and trading speeds when each asset's trading cost is also state-dependent. In the second experiment, we analyze the sensitivity of the risk-parity allocation strategy to stochastic trading costs.

Table 1: Parameter Values for Numerical Experiments 1 and 2. This table reports the parameter values used in the numerical experiments described in Section 3.1 and 3.2. Trading is daily, and reported values of μ and Σ are annualized.

<i>Experiment 1 (Corporate Bonds)</i>		<i>Experiment 2 (Risk-parity)</i>	
Parameter	Value	Parameter	Value
γ	10^{-8}	γ	10^{-8}
ρ	0.9996	ρ	0.9996
π_{LL}	0.95	π_{LL}	0.95
π_{HH}	0.9	π_{HH}	0.9
μ_L	$\begin{bmatrix} 10 \\ 8 \end{bmatrix}$	μ_L	$\begin{bmatrix} 1 \\ 1 \end{bmatrix}$
μ_H	$\begin{bmatrix} 12 \\ 16 \end{bmatrix}$	μ_H	$\begin{bmatrix} 1 \\ 1 \end{bmatrix}$
Σ_L	$\begin{bmatrix} 100 & 50 \\ 50 & 100 \end{bmatrix}$	Σ_L	$\begin{bmatrix} 100 & 0 \\ 0 & 900 \end{bmatrix}$
Σ_H	$\begin{bmatrix} 900 & 450 \\ 450 & 900 \end{bmatrix}$	Σ_H	$\begin{bmatrix} 400 & 0 \\ 0 & 3600 \end{bmatrix}$
Λ_L	$\begin{bmatrix} 1.25 \times 10^{-8} & 0 \\ 0 & 10^{-8} \end{bmatrix}$	Λ_L	$5 \times 10^{-8} \Sigma_L$
Λ_H	$\begin{bmatrix} \text{Variable} & 0 \\ 0 & 10^{-8} \end{bmatrix}$	Λ_H	(Variable) $\eta \Sigma_H$

3.1 Corporate Bonds vs. Treasuries

To illustrate the implications of the model, we consider a case with two assets and two states: Low-risk (L) and High-risk (H). In the low-risk state, Asset 1 (e.g., “Corporate”) has higher Sharpe ratio than Asset 2 (e.g., “Treasury”). However, in the high-risk state, Asset 2 has higher Sharpe ratio. In both states, Asset 2 is cheaper to trade than Asset 1. We assume that both assets are positively correlated.

Table 1 shows a simple calibration for this example. We assume that the initial price for the assets are \$100, e.g., Asset 1 has an annual expected price change of \$10 in the L -state and its volatility is \$10 in this state. Trading frequency is daily.

In the right panel of Figure 1 we plot the conditional Markowitz portfolios. In the left panel, we plot the aim portfolios in both states as we vary the transaction costs of the first asset in the high-risk state. The minimum value is 1.25×10^{-8} and can go up to 2×10^{-7} . We see that the aim portfolio in the high state is always very close to the Markowitz portfolio. Intuitively, since price impact is very small in the low-risk state, the aim portfolio in the high-risk state needs not take into account the investment opportunity set in the low-risk state. Instead, in the low-risk state, as we increase price impact of Asset 1 in the high-risk state, the aim portfolio varies dramatically. It weighs more and more heavily the Markowitz portfolio in the high-risk state, thus lowering the

Figure 1: Aim and Conditional Markowitz Portfolios for Numerical Experiment 1

desired position in Asset 1 and increasing the desired position in Asset 2. Eventually, for high enough expected trading costs of Asset 1 in the H -state it becomes optimal to hold more of Asset 2 even in the low-risk state. That is, it is optimal to hold more of the asset that appears dominated in Sharpe ratio terms in the L -state preemptively to anticipate the future desired deleveraging in the H -state.

Figure 2 plots the corresponding trading speeds in both assets in both regimes.¹⁴ Intuitively, we see that the trading speed is generally higher in the high-risk regime due to the higher tracking error cost. However, as it becomes more costly to trade Asset 1 in that regime, its trading speed drops rapidly. Interestingly, the trading speed of Asset 1 actually increases in the Low-risk regime in response to the increase of its trading cost in the high-risk regime. That is even though Asset 1 has a marginally higher trading cost than Asset 2 in the low-risk regime, it is optimal to trade it more aggressively than Asset 2 in the low-risk regime in anticipation of its much higher trading cost in the high-risk regime.

This example captures some salient features of the Corporate versus Treasury bond returns. Like Asset 1, Corporate bonds typically offer higher expected rates of returns in expansions (good states) than Treasury bonds (Asset 2). However, during recessions (bad states) their risk increases dramatically and the higher probability of default leads to lower returns.¹⁵ Further, it is also a fact that corporate bonds become a lot costlier to trade in bad states than Treasuries, whose liquidity remains very high. As the stylized example demonstrates, because it is optimal to reduce the position in the Corporates in the high-risk state when these are very costly to trade, it can be optimal to hold a larger share of the Treasuries already in the good state even though in that state

¹⁴For simplicity, we only plot the diagonal values of the trading speed matrix $\Lambda_s^{-1}Q_s$, which is actually not diagonal in this example.

¹⁵Of course, it is arguable whether the *expected return* is actually lower, since expected returns are hard to measure. For illustration we assume that in the bad states the risk of Asset 1 is higher and its Sharpe ratio is lower than that of Asset 2.

Figure 2: Trading Speed for numerical experiment 1

the conditional Sharpe ratio of Corporates dominates that of Treasuries. Further, even though Corporates may be more costly to trade than Treasuries in the good state, it may be optimal to trade them more aggressively in the good state in anticipation of the high-risk regime with much higher relative trading costs.

This example helps also to think about the question: in a portfolio with liquid and illiquid assets, which one should one liquidate first because of a liquidity shock? Our analysis gives the following answer. First and foremost, one should trade the illiquid asset more aggressively in anticipation of the future liquidity crisis and steer the portfolio to a position that overweights liquid assets, possibly deviating from the unconditional optimal portfolio to take into account the future possible risk and liquidity shocks. Second, once the crisis hits, one should trade less aggressively the more costly assets and more aggressively the liquid assets to steer the portfolio towards the conditional mean-variance efficient portfolio.

3.2 Risk-parity Allocation Strategy

The risk-parity allocation strategy has received a lot of attention among practitioners, not the least because it is applied in the very successful “All-weather” fund of Bridgewater. As a rationale for such a strategy, it is sometimes argued that such a strategy reflects the difficulty in measuring expected returns of and correlations between asset classes. With all expected price changes equal and constant (e.g., $\mu = \mathbf{1}$) and all correlation coefficients equal to zero, the mean-variance efficient Markowitz portfolio becomes a risk-parity portfolio in the sense that every asset contributes an identical amount of volatility to the overall portfolio ($m_s = (\gamma \Sigma_s)^{-1} \mathbf{1} = \text{diag}(\frac{1}{\gamma v_{i,s}})$).¹⁶ Here we illustrate that it is actually optimal to deviate from the risk-parity allocation under these same assumptions, if transaction costs of various asset classes move predictably with their volatility.

¹⁶See (Asness, Frazzini, and Pedersen 2012) for further discussion.

Figure 3: Aim portfolios for numerical experiment 2

With $\Sigma_s = \text{diag}(v_{i,s})$ and $\Lambda_s = \text{diag}(\lambda_{i,s})$ and $\mu_s = \mathbf{1}$, we can solve for the optimal aim portfolio in closed-form from Lemma 3 with $F = \text{diag}(\mathbf{1})$.

We illustrate in Figure 3 how the aim portfolio in state s can deviate from the equal risk portfolio as transaction costs in state s' increase. Comparing the left panel to the right panel, we see that in the high-risk state, the aim portfolio remains very close to the risk-parity portfolio. This is because the risk-parity portfolio is the conditional Markowitz portfolio in state s under our assumptions and we need not take into account the low-risk state, since rebalancing is expected to be much less costly then. However, in the low-risk state, it is optimal to deviate dramatically from the equal-risk weights. Indeed, we need to lower considerably the target position in both assets, but especially in the high-risk asset, in anticipation of the future increase in volatility and in the cost of trading that asset, in case of a switch to the high-risk state.

Trading speeds for all assets are plotted in Figure 4. As we can see, trading speed decreases in the high-risk state and increases in the low-risk state when transaction costs in the high-risk state are increasing. That is the more costly it becomes to trade assets in the H -state, the more aggressively we have to trade assets in the low-risk state. We note that trading speeds are not security specific in this experiment, because we assume that the price impact matrix is a constant multiple of the covariance matrix.

4 A Regime Switching Model for Returns

Following much of the literature (e.g., GP, Litterman), the model in the earlier section assumes that conditional on a state, the expectation and covariance matrix of price changes are constant. This leads to a very tractable solution, because in a mean-variance framework the only motive to rebalance the portfolio conditional on holding the mean-variance efficient portfolio, is if there is a change in the state, that is if there is a change in the expectation or covariance matrix of

Figure 4: Trading Speeds for numerical experiment 2

price changes. Unfortunately, it is not a very plausible model for returns empirically, as it assumes counterfactual dynamics for the return covariances. Mei, DeMiguel, and Nogales (2016) document that the assumption of stationary price changes is reasonable with short-term horizons of up to one year. Empirically, the “conditional log-normal” model of price changes is preferable to the “conditional normal” model assumed in this section. Interestingly, in our framework the “log-normal” model, which assumes that the expectation and covariance matrix of dollar returns is constant in a given state, is very tractable as well.

In this section, we present a regime switching model formulated in dollars and returns as opposed to price changes and number of shares. In our empirical analysis in Section 5 we apply this model to timing the market portfolio while accounting for time-varying transaction costs and stochastic volatility.

4.1 Formulation

We have N risky assets and collect the N -dimensional vector of returns from period t to $t + 1$ in $r_{t+1} \equiv \frac{dS_t}{S_t}$. The net return vector has the following state-dependent mean and covariances:

$$\begin{aligned} E[r_{t+1}] &= \mu(s_t) \\ E[(r_{t+1} - \mu(s_t))(r_{t+1} - \mu(s_t))^T] &= \Sigma(s_t) \end{aligned}$$

$\mu(s_t)$ and $\Sigma(s_t)$ are, respectively, the N -vector of expected returns and the $N \times N$ covariance matrix of returns. Both μ and Σ are a function of a state variable s_t which follows a Markov chain with transition probabilities $\pi_{s,s'}$.

Since the model is set-up in dollars, the investor rebalances at the end of each period again in dollars. If the dollar trade vector is given by u_t , then, the dollar holdings of the investor has the

following dynamics:

$$x_{t+1} = \text{diag}(\mathbf{1} + r_{t+1})x_t + u_{t+1} \quad (8)$$

$$= \text{diag}(R_{t+1})x_t + u_{t+1} \quad (9)$$

where the gross returns are given by R_{t+1} .

We consider the optimization problem of an agent with the following objective function with an infinite investment horizon:¹⁷

$$\max_{x_t} \mathbb{E} \left[\sum_{t=1}^{\infty} \rho^{t-1} \left\{ x_t^\top \mu(s_t) - \frac{1}{2} \gamma x_t^\top \Sigma(s_t) x_t - \frac{1}{2} u_t^\top \Lambda(s_t) u_t \right\} \right] \quad (10)$$

The agent chooses her dollar holdings x_t in each period t so as to maximize this objective function. Specifically, at the end of period $t-1$, the agent holds x_{t-1} dollars. At this point the agent observes the state s_t , and trades u_t dollars to bring his dollar holdings to $\text{diag}(R_t)x_t + u_t$. We again consider a linear price impact model. The total (dollar) cost of trading u_t is $\frac{1}{2} u_t^\top \Lambda(s_t) u_t$.

4.2 Value Functions and Optimal Portfolio

For simplicity, we consider a two-state Markov chain model, with states H and L . The model is straightforward to generalize to multiple states. In our empirical application in Section 5 we consider two-state and four-state models. Using the dynamic programming principle, the value function $V(x_{t-1}, R_t, s_t)$ satisfies

$$V(x_{t-1}, R_t, s) = \max_{x_t} \left(x_t^\top \mu_s - \frac{1}{2} u_t^\top \Lambda_s u_t - \frac{\gamma}{2} x_t^\top \Sigma_s x_t + \rho \mathbb{E}_t [V(x_t, \mathbf{1} + \mu_s + \epsilon_s, z)] \right),$$

where $\mathbb{E}[\epsilon_s] = 0$ and $\mathbb{E}[\epsilon_s \epsilon_s^\top] = \Sigma_s$. We guess the following quadratic form for our value functions:

$$V(x, R, s) = -\frac{1}{2} x^\top \text{diag}(R) Q_s \text{diag}(R) x + x^\top \text{diag}(R) q_s + c_s,$$

where Q_s is a symmetric $N \times N$ matrix and q_s, c_s are N dimensional vectors of constants for $s \in \{H, L\}$. We can now simplify $\mathbb{E}_t [V(x_t, \mathbf{1} + \mu_s + \epsilon_s, z)]$ using the assumed structure for the value functions and write it in the form of $-\frac{1}{2} x_t^\top A_s x_t + x_t^\top b_s + d_s$ where

$$Z_s = \mathbb{E}[(\mathbf{1} + \mu_s + \epsilon_s)(\mathbf{1} + \mu_s + \epsilon_s)^\top] = \Sigma_s + (\mathbf{1} + \mu_s)(\mathbf{1} + \mu_s)^\top,$$

$$A_s = \pi_{s,s}(Z_s \circ Q_s) + \pi_{s,s'}(Z_s \circ Q_{s'}),$$

$$b_s = \pi_{s,s}(\mu_s \circ q_s) + \pi_{s,s'}(\mu_s \circ q_{s'}),$$

$$d_s = \pi_{s,s}c_s + \pi_{s,s'}c_{s'},$$

¹⁷In the appendix, as discussed in footnote 9 we provide two ways to micro-found this objective function.

and $\circ$ denotes element-wise multiplication. Using this expression for $\mathbb{E}_t[V(x_t, \mathbf{1} + \mu_s + \epsilon_s, z)]$, we obtain

$$V(x_{t-1}, R_t, s) = \max_{x_t} \left\{ x_t^\top \mu_s - \frac{1}{2} (x_t - \text{diag}(R_t)x_{t-1})^\top \Lambda_s (x_t - \text{diag}(R_t)x_{t-1}) - \frac{\gamma}{2} x_t^\top \Sigma_s x_t - \frac{\rho}{2} x_t^\top A_s x_t + \rho x_t^\top b_s + \rho d_s \right\},$$

Thus, we maximize the quadratic objective $-\frac{1}{2}x_t^\top J_s x_t + x_t^\top j_s + k_s$ where we define

$$\begin{aligned} J_s &= \gamma \Sigma_s + \Lambda_s + \rho A_s \\ j_s &= \Lambda_s \text{diag}(R_t)x_{t-1} + \mu^s + \rho b_s \\ k_s &= -\frac{1}{2}x_{t-1}^\top \text{diag}(R_t)\Lambda_s \text{diag}(R_t)x_{t-1} + \rho d_s \end{aligned}$$

Then, the optimal x_t when the state is s is given by $J_s^{-1}j_s$. That is to say

$$x_t = (\gamma \Sigma_s + \Lambda_s + \rho A_s)^{-1} (\Lambda_s \text{diag}(R_t)x_{t-1} + \mu^s + \rho b_s). \quad (11)$$

The value achieved at the optimal solution is given by $\frac{1}{2}j_s^\top J_s^{-1}j_s + k_s$ and we obtain the following coupled matrix equations:

$$Q_s = -\Lambda_s (\gamma \Sigma_s + \Lambda_s + \rho A_s)^{-1} \Lambda_s + \Lambda_s, \quad (12)$$

$$q_s = \Lambda_s (\gamma \Sigma_s + \Lambda_s + \rho A_s)^{-1} (\mu_s + \rho b_s), \quad (13)$$

$$c_s = \frac{1}{2}(\mu_s + \rho b_s)^\top (\gamma \Sigma_s + \Lambda_s + \rho A_s)^{-1} (\mu_s + \rho b_s) + \rho d_s. \quad (14)$$

Overall, these equations are very similar to those obtained in the previous section for the regime switching model of price changes. The main difference is the need to introduce the matrices A_s and b_s which are non-linear transformations of Q_s and q_s . We solve for Q_s and q_s iteratively from equations (12) and (13) respectively. We use the zero matrix for Q_s and the zero vector for q_s as initial guesses. Convergence is obtained very rapidly in all of our implementations.

4.3 Aim Portfolio and Trading Speed

Following our analysis in the previous section, we define the aim portfolio in each state, aim_s , as the portfolio at which it would be optimal not to rebalance given the current state s . The following lemma characterizes the aim portfolio and the trading speed.

Lemma 3 *The conditional aim portfolio aim_s at which it is optimal not to rebalance is given by*

$$aim_s = (\gamma \Sigma_s + \rho A_s)^{-1} (\mu_s + \rho b_s)$$

It maximizes the value function $V(x_{t-1}, R_t, s)$ with respect to $x_{t-1} \text{diag}(R_t)$.

The optimal trading rule is to “trade partially towards the aim” at the trading speed $\tau_s = \Lambda_s^{-1} Q_s$:

$$x_s = (I - \tau_s) \text{diag}(R_t) x_{t-1} + \tau_s \text{aim}_s$$

Proof. Maximizing the value function at time $V(x_{t-1}, R_t, s)$ with respect to $\text{diag}(R_t) x_{t-1}$ we obtain:

$$\text{aim}_s = Q_s^{-1} q_s$$

Substituting from the definitions in equations (12) and (13) we obtain:

$$\begin{aligned} \text{aim}_s &= \left(-\Lambda_s (\gamma \Sigma_s + \Lambda_s + \rho A_s)^{-1} \Lambda_s + \Lambda_s \right)^{-1} \left(\Lambda_s (\gamma \Sigma_s + \Lambda_s + \rho A_s)^{-1} (\mu_s + \rho b_s) \right) \\ &= \left(-(\gamma \Sigma_s + \Lambda_s + \rho A_s)^{-1} \Lambda_s + I \right)^{-1} (\gamma \Sigma_s + \Lambda_s + \rho A_s)^{-1} (\mu_s + \rho b_s) \\ &= (\gamma \Sigma_s + \rho A_s)^{-1} (\mu_s + \rho b_s) \end{aligned}$$

where the last equality obtains by noting that if we define the matrix

$$M = \left(-(\gamma \Sigma_s + \Lambda_s + \rho A_s)^{-1} \Lambda_s + I \right)^{-1} (\gamma \Sigma_s + \Lambda_s + \rho A_s)^{-1}$$

then

$$M^{-1} = (\gamma \Sigma_s + \Lambda_s + \rho A_s) \left(-(\gamma \Sigma_s + \Lambda_s + \rho A_s)^{-1} \Lambda_s + I \right) = (\gamma \Sigma_s + \rho A_s),$$

which immediately implies that $M = (\gamma \Sigma_s + \rho A_s)^{-1}$.

To prove the second part of the lemma, we start from the definition of the optimal position x_t given in equation (11). It is straightforward to obtain the optimal trade

$$x_t - \text{diag}(R_t) x_{t-1} = \Lambda_s^{-1} (\gamma \Sigma_s + \rho A_s) (\text{aim}_s - x_t).$$

Using the definition of matrix M above and equation (12), we obtain the formula for the trading speed. ■

4.4 Difference between Two Models

Figure 5 compares aim portfolios and trading speeds in models set-up in shares and dollars. We calibrate the model to a share price worth one dollar so that the y -axis represents the dollar investment of both strategies (that is number of shares invested equal number of dollars invested). Table 2 displays all of the model parameters. We see that the aim portfolio in the regime switching model of price changes always invests a larger position in the risky asset than the aim portfolio for the regime switching model of returns. The difference is larger the larger the expected return on the stock. The intuition is that when we rebalance at time t in the return model, the dollar position will be affected by the risky return (see equation (8)), before we get to rebalance. Thus, the aim portfolio in dollars reflects the expected dollar position after the risky one period return is

Figure 5: Aim portfolios and trading speeds in models set-up in shares and dollars.

realized.

We also observe that the trading speed is higher for the regime-switching-model of returns than for that of price changes. This is because, in the regime-switching model of returns, there is an additional “rebalancing motive” for trading, as dollar positions drift away from their target as a result of return shocks (even in the absence of any change in the investment opportunity set).

5 Empirical Application

In this section, we implement our methodology using the modeling framework in dollars and illustrate that there are economically significant benefits using our approach both in-sample and out-of-sample.

5.1 Model Calibration

We use daily value weighted CRSP market returns from 1967 Q3 to 2017 Q2 (50 years) to estimate a regime switching model. The data is downloaded from Ken French’s data library.

Guidolin and Timmermann (2006) consider a range of values for the number of states and find that a four-state regime model performs better in explaining bond and stock returns. Following this study, we estimate a Markov switching model with four states to describe the dynamics of market returns:

$$r_{t+1} = \mu(s_t) + \sigma(s_t)\epsilon_{t+1} \quad (15)$$

where $s_t = \{1, 2, 3, 4\}$ and ϵ_{t+1} are serially independent and drawn from standard normal distribution. State transitions occur according to a Markov chain and we denote by P_{ij} the probability of

Table 2: Parameter Values	
Parameter	Value
γ	5×10^{-8}
ρ	0.9996
π_{LL}	0.98
π_{HH}	0.9
μ_L	[0]
μ_H	(Variable)
Σ_L	$[0.40 \times 10^{-4}]$
Σ_H	$[3.33 \times 10^{-4}]$
Λ_L	$[2 \times 10^{-10}]$
Λ_H	$[3 \times 10^{-10}]$

switching from state i to state j .¹⁸

Table 3 displays the estimates of the model. All coefficients are statistically significant at 1% level. Overall, we observe that the rank correlation between the estimated expected returns and volatilities is not equal to 1. We observe that the expected return can be lower in a high volatility state. This pattern has been found since the initial applications with regime switches on equity returns (see e.g., Hamilton and Susmel (1994)).

The top two panels in Figure 6 illustrate the corresponding smoothed probabilities for each regime and the bottom panel in Figure 6 illustrates the color-coded regimes by using the maximum smoothed probability for identification. The first regime (green) highlights the good states of the return data with high return and low volatility corresponding to the highest Sharpe ratio. This regime has also the highest expected duration with roughly 51 trading days. The transition from this state usually occurs to the second state (blue) with slightly lower expected return and higher volatility. The expected duration for this state is 30 trading days. The third state (yellow) is a distressed state with low expected return and high volatility. This state has the lowest Sharpe ratio and has an expected duration of approximately 33 trading days. The final state covers the crisis periods with very high expected return and very high volatility. We observe that it covers trading days around the 1987 crash, the dot-com bubble and the financial crisis. This state is relatively short-lived with an expected duration of roughly 15 trading days.

¹⁸To restrict the number of parameters, we have also tried fitting a four-state model that constrains the general model to having only two mean and volatility coefficients (i.e., mean or volatility may remain unchanged after a transition) as opposed to four but this constrained model can be rejected with a likelihood test.

Table 3: Parameter estimates for a four-state regime switching model using daily market return data from 1967 Q3 and 2017 Q2.

Parameter	Estimate	Parameter	Estimate
μ_1	0.0864%	σ_1	0.5512%
μ_2	0.0340%	σ_2	0.9372%
μ_3	0.0069%	σ_3	1.6032%
μ_4	0.2939%	σ_4	3.9178%
P_{11}	0.9804	P_{12}	0.0196
P_{13}	0.0000	P_{14}	0.0000
P_{21}	0.0250	P_{22}	0.9670
P_{23}	0.0080	P_{24}	0.0000
P_{31}	0.0000	P_{32}	0.0233
P_{33}	0.9693	P_{34}	0.0074
P_{41}	0.0016	P_{42}	0.0000
P_{43}	0.0635	P_{44}	0.9350

5.2 Calibration of the Transaction Costs

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

The execution data covers S&P 500 stocks between January 2011 and December 2012. Execution duration is greater than 5 minutes but no longer than a full trading day. Total number of orders is 81,744 with an average size of approximately \$1 million. The average participation rate of the order, the ratio of the order size to the total volume realized in the market, is approximately 6%. Table 4 reports further summary statistics on the large-order execution data.

According to our quadratic transaction cost model, trading q dollars in state j would cost the investor $\lambda_j q^2$. Since each of our executions are completed in a day, we can uniquely label each execution originating in one of the four states by setting it to the state with maximal smoothed probability. With this methodology, we find that 22,946 executions are in regime 1, 41,898 execu-

Table 4: Summary statistics for the main attributes in the execution data. Participation rate is equal to the ratio of the executed volume to total volume during the lifetime of the order. The volatility of the asset is estimated using the mid-quote prices. Order duration is expressed as a fraction of full trading day (i.e., 6.5 hours).

Statistic	Mean	Min	Pctl(25)	Median	Pctl(75)	Max
Order Value (\$ M)	0.967	0.001	0.094	0.396	1.102	158.300
Participation Rate	0.061	0.00001	0.002	0.013	0.102	1.000
Volatility	0.014	0.0002	0.008	0.011	0.016	0.344
Order Duration	0.384	0.013	0.041	0.153	0.851	1.000
IS (bps)	4.095	-1006.000	-16.440	4.075	25.080	996.600

tions in regime 2, 14,502 executions in regime 3 and 2,398 in regime 4. Compared to other states, regime 4 has relatively small number of executions due to its short-lived nature. At first sight, it is surprising that we have the largest number of executions in regime 2. But, during the 2011-2012 period, the volatility was relatively high so there are actually fewer trading days in regime 1 (see Figure 6).

Our execution data has information on both the order size and total trading cost. Total trading cost is computed by comparing the average price of the execution to the prevailing price in the market before the execution starts. This is usually referred to as implementation shortfall (*IS*) (Perold 1988). Formally, *IS* of the i th execution is given by

$$IS_i = \text{sgn}(Q_i) \frac{P_i^{\text{avg}} - P_{i,0}}{P_{i,0}}, \quad (16)$$

where Q_i is the dollar size of the order (negative if a sell order), P_i^{avg} is the volume-weighted execution price of the parent-order and $P_{i,0}$ is the average of the bid and ask price at the start-time of the execution. Thus, total trading cost in dollars is equal to $IS_i \times Q_i$. According to our model, this is given by $\lambda_{m(i)} Q_i^2$ where $m(i)$ maps the i th execution to the state of the trading day. Thus, we can estimate λ_j for each state by fitting the following model:

$$IS_i = \lambda_1 Q_i \mathbf{1}_{\{m(i)=1\}} + \lambda_2 Q_i \mathbf{1}_{\{m(i)=2\}} + \lambda_3 Q_i \mathbf{1}_{\{m(i)=3\}} + \lambda_4 Q_i \mathbf{1}_{\{m(i)=4\}} + \varepsilon_i$$

Table 5 illustrates the estimated coefficients. The reported standard errors are clustered at the stock and calendar day level. We observe that λ estimates are all highly significant (except in state 4 where we observe fewer executions in our data-set) and vary a lot across regimes and tend to increase with volatility. We find that λ_3 is the largest across all states. Recall that this state has the lowest Sharpe ratio and thus can be interpreted as the distressed state. Using Wald tests pairwise, we find that the estimate of transaction costs in this distressed state, λ_3 , is statistically higher than all other coefficients at a 10% significance level.

To better understand the variation in transaction costs across our states, we present in Table 6 the average values of various liquidity proxies in each state. We find that bid-ask spreads, mid-quote

Table 5: Transaction cost estimates in each regime labeled from the four-state regime switching model. λ_n denotes the transaction cost multiplier in regime n . The second column reports the results from a liquid subset in which we only include executions from stocks within the top 10% in market capitalization. Estimated values are multiplied by 10^{10} . Standard errors are double-clustered at the stock and calendar day level.

	<i>Dependent variable: IS</i>	
	<i>All Stocks</i>	<i>Liquid</i>
λ_1	1.688*** (0.459)	0.501** (0.217)
λ_2	1.725*** (0.195)	0.793*** (0.189)
λ_3	3.037*** (0.418)	1.506*** (0.352)
λ_4	2.274 (1.927)	0.935 (1.329)
<i>Note:</i> *p<0.1; **p<0.05; ***p<0.01		

volatility and turnover are increasing across states, i.e., volatility. However, the Amihud illiquidity proxy returns similar ranking to the estimated λ coefficients with state 3 being more illiquid than state 4. Since volume is much larger in that state, it may act as a mitigating factor on trading costs (see e.g., Admati and Pfleiderer (1988) and Foster and Viswanathan (1993)).

Since we would like to estimate the price impact of trading the market portfolio, our estimates may be overestimating the cost as it is based on the complete set of S&P 500 stocks. In order to address this issue, we rerun our regressions only using data corresponding to the top 10% of stocks with respect to market capitalization. We believe that this universe of stocks reflect a more natural comparison to the market portfolio.

The second column of Table 5 illustrates the estimated coefficients for this liquid universe. We observe that the coefficients are lower by a factor between two and three but preserve the same ranking across states. In this case, λ_3 is statistically different than the coefficients of the first and second state at 10% significance level. The second panel of Table 6 illustrates the average values of each liquidity proxy in each regime using this universe of large-cap stocks.

5.3 Objective function

We use the regime switching model based on dollar holdings and returns presented in Section 4 as the investment horizon is very long. Formally, the investor's objective function is:

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \rho^t \left\{ x_t \mu(s_t) - \frac{1}{2} \lambda(s_t) u_t^2 - \frac{\gamma}{2} \sigma^2(s_t) x_t^2 \right\} \right] \quad (17)$$

Table 6: Average liquidity proxies in each regime. The second column reports the averages from a liquid subset in which we only include executions from stocks within the top 10% in market capitalization. Standard errors are double-clustered at the stock and calendar day level.

	All Stocks				Liquid			
	Spread (bps)	Volatility (%)	Turnover	Amihud ($\times 10^{-8}$)	Spread (bps)	Volatility (%)	Turnover	Amihud ($\times 10^{-8}$)
1	3.80*** (0.04)	1.11*** (0.02)	3.51*** (0.12)	1.37*** (0.06)	2.89*** (0.03)	1.01*** (0.02)	2.93*** (0.11)	0.47*** (0.02)
2	3.95*** (0.04)	1.23*** (0.02)	3.58*** (0.13)	1.59*** (0.06)	2.98*** (0.03)	1.13*** (0.02)	3.02*** (0.10)	0.55*** (0.02)
3	4.95*** (0.09)	1.92*** (0.06)	4.27*** (0.20)	1.97*** (0.08)	3.52*** (0.07)	1.77*** (0.06)	3.63*** (0.19)	0.67*** (0.04)
4	5.62*** (0.39)	2.84*** (0.28)	7.37*** (2.07)	1.80*** (0.34)	3.97*** (0.27)	2.59*** (0.25)	6.42*** (1.07)	0.59*** (0.12)

Note: *p<0.1; **p<0.05; ***p<0.01

where $x_t = x_{t-1}(1 + r_t) + u_t$ and $s_t \in \{1, 2, 3, 4\}$. We calibrate ρ so that the annualized discount rate is 1%. We set $\gamma = 1 \times 10^{-10}$ which we can think of as corresponding to a relative risk aversion of 1 for an agent with \$10 billion dollars under management. We assume that the investor starts from zero holdings and rebalances daily.

The optimal portfolio policy of the investor is given by

$$x_t^{\text{opt}}(s_t) = (1 - \frac{Q(s_t)}{\lambda(s_t)})(1 + r_t)x_{t-1}^{\text{opt}} + \frac{Q(s_t)}{\lambda(s_t)}aim(s_t) \quad \forall s_t \in \{1, 2, 3, 4\} \quad (18)$$

where q and Q solve the following system of equations $\forall s \in \{1, 2, 3, 4\}$:

$$Q(s_t) = -\lambda(s_t)^2 (\gamma\sigma^2(s_t) + \lambda(s_t) + \rho(\sigma^2(s_t) + (1 + \mu(s_t))^2) \bar{Q}(s_t))^{-1} + \lambda(s_t), \quad (19)$$

$$q(s_t) = (\mu(s_t) + \rho\mu(s_t)\bar{q}(s_t)) \left(1 - \frac{Q(s_t)}{\lambda(s_t)}\right), \quad (20)$$

$$aim(s_t) = Q(s_t)^{-1}q(s_t). \quad (21)$$

Since we have only one asset, the trading speed is one-dimensional and given by $\frac{Q(s_t)}{\lambda(s_t)}$ in each state.

5.4 Aim Portfolios and Trading Speed

Using the estimated model coefficients, we first study the aim portfolios across states in the presence and absence of transaction costs. Figure 7 illustrates the aim portfolios for the optimal policy in these cases. We also compare this optimal policy with a simple unconditional mean-variance benchmark, in which the portfolio rule holds a constant dollar amount equal to $\frac{\mu_{\text{avg}}}{\gamma\sigma_{\text{avg}}^2}$ in the risky

asset. Here, μ_{avg} and σ_{avg}^2 are the sample mean and variance of the market returns between 1967 Q3 and 2017 Q2.

In the top panel, the red solid line illustrates the aim portfolios in the absence of transaction costs. Without transaction costs, aim portfolios are simply the conditional mean-variance optimal Markowitz portfolios. Compared to the unconditional mean-variance constant benchmark portfolio, the conditional Markowitz portfolio is very aggressive in regime 1 and holds a smaller amount than the constant portfolio in all other states. In regime 3, the holdings are very close to a risk-free position.

In the bottom panel, we plot the aim portfolios when there are stochastic trading costs. We use the estimated transaction cost multipliers from the liquid subset as provided in Table 5. Surprisingly, regime 4 has the smallest aim portfolio whereas regime 3, the lowest Sharpe ratio state, has slightly higher holdings. This is due to differences in trading costs, as well as to the transition probabilities, across states. For example, trading costs are largest in Regime 3, thus the optimal aim portfolio, which will determine trading in that state, should depend on the average positions expected in states that it will transition from, essentially Regime 4 (probability of $\approx 6\%$) and Regime 2 (probability of $\approx 1\%$), as well as from states it will transition too, again Regime 2 (probability of $\approx 2\%$) and Regime 4 (probability of $\approx 1\%$). These considerations make the desired holdings in Regime 3 higher. Interestingly, the aim portfolios in regime 3 and regime 4 hold a larger position in risky assets than the corresponding conditional Markowitz portfolios, whereas the aim portfolio in regime 1 actually holds a much smaller position than the conditional Markowitz portfolio. This emphasizes the impact of transaction costs and potential transitions between states on desired holdings.

Finally, we plot the trading speeds in each regime in Figure 8. Due to high volatility, regime 4 has the highest trading speed. Regime 1 has the lowest trading costs so we find that the trading speed is relatively larger compared to regime 2 and regime 3. However, the difference is not very large as these other states have higher volatilities. Regime 3 has the lowest trading speed potentially due to its highest trading costs.

5.5 In-sample Analysis

In this section, we evaluate the performance of the optimal policy using the in-sample estimates from our four-state regime switching model. We compare it to various benchmark policies in the presence and absence of transaction costs to quantify the potential benefits of this methodology.

In order to evaluate the performance of the policies, we need to assign each trading day to a regime state so that we can determine the appropriate values of $\sigma^2(s_t)$ and $\lambda(s_t)$. For this purpose, we use the smoothed probabilities from the regime switching model and assign the regime of each trading day to the state with the highest smoothed probability.

We also skip a day to implement the optimal and myopic policies without any forward-looking bias. That is to say, to determine the position on day t , we use the smoothed probabilities from day $t - 1$.

Let x_t^{opt} be the optimal policy as computed from Equation (18) and the above implementation methodology. We break down the realized objective function into two terms, wealth and risk penalties:

$$W_T^{\text{opt}} = \sum_{t=1}^{T=12600} \rho^t \left[x_t^{\text{opt}} r_{t+1} - \frac{1}{2} \lambda(s_t) \left(x_t^{\text{opt}} - x_{t-1}^{\text{opt}} R_t \right)^2 \right] \quad (22)$$

$$RP_T^{\text{opt}} = \sum_{t=1}^{T=12600} \frac{1}{2} \rho^t \left[x_t^{\text{opt}} \sigma(s_t) \right]^2 \quad (23)$$

Here, $t = 12600$ corresponds to the final trading day of 2017 Q2.

5.5.1 Benchmark Policies

As described earlier, the first benchmark policy is the constant-dollar rule in which the investor chooses $x_t^{\text{con}} = \frac{c\mu_{\text{avg}}}{\gamma\sigma_{\text{avg}}^2}$. The parameters, μ_{avg} and σ_{avg}^2 , are obtained using the full in-sample data. We choose c so that the policy has the same risk exposure as the optimal policy, i.e., the discounted sum of risk penalties from this policy equals RP_T^{opt} . In the presence of trading costs, getting into a large constant position in the first period may result in large trading costs so to minimize this effect we allow this policy to build the constant position in the first 10 trading days with equal-sized trades.

The second benchmark policy is the buy-and-hold portfolio in which the investor invests x_0 dollars into the market portfolio at the beginning of the horizon.¹⁹ We provide slight advantage to this benchmark policy by assuming that he builds this position with no trading costs. The investor never trades till the end of the investment horizon. We again optimally choose x_0 so that the policy has the same risk exposure as the optimal policy.

The third benchmark policy is the myopic policy with transaction cost multiplier, a widely used practitioner approach. This approach solves a myopic mean-variance problem, that is given some initial position (x_{t-1}) and the state s_t, r_t , it solves $\max_{u_t} x_t \mu(s_t) - \frac{1}{2} \gamma \sigma(s_t)^2 x_t^2 - \frac{1}{2} h u_t^2 \lambda(s_t)$ subject to the dynamics $x_t = x_{t-1}(1 + r_t) + u_t$. The myopic policy with transaction cost multiplier h is given by

$$x_t^{\text{my}}(s_t) = (1 - \tau(s_t))(1 + r_t)x_{t-1}^{\text{my}} + \tau(s_t) \frac{\mu(s_t)}{\gamma\sigma^2(s_t)} \quad \forall s_t \in \{1, 2, 3, 4\} \quad (24)$$

$$\tau(s_t) = \frac{1}{1 + \frac{h\lambda(s_t)}{\gamma\sigma^2(s_t)}} \quad (25)$$

Note that this policy, like the optimal one, trades partially towards an aim portfolio. However, since it takes the current state as given and ignores the implications of any future transitions in the state, the aim portfolio is the conditional mean-variance efficient Markowitz portfolio and the

¹⁹We assume that the investor shorts the risk-free asset to generate this initial capital so he also starts from zero wealth.

trading inertia, $1 - \tau(s_t) \approx \frac{h\lambda(s_t)}{\gamma\sigma^2(s_t)}$, only depends on the ratio between current state's transaction costs and the variance. We choose h so that the myopic policy uses the optimal trading speed $\tau^*(s_t)$ in each regime. Note that in this case, the risk penalties will not be the same. Further, in the absence of transaction costs, the myopic policy is optimal, thus, we compare it to the optimal one only in the presence of transaction costs.

5.5.2 Comparison between Portfolio Policies

Figure 9 compares the optimal policy to the constant portfolio in the absence of trading costs. Both policies have the same risk penalty by construction (see bottom-right panel), thus the wealth dynamics are direct measures of performance. The top-left panel illustrates that the optimal policy has a much higher performance. We observe that this is achieved by trading more and timing the regimes of the return data. This confirms that there is predictability and that, at least in the absence of transaction costs, there is value to rebalancing across the estimated regimes.

Figure 10 compares the optimal policy to the buy-and-hold portfolio in the absence of trading costs. Both policies again have the same risk penalty by construction. In the top-right panel, the starting position for the buy-and-hold policy is roughly $\$3 \times 10^9$. Since this policy never trades, the position becomes very large at the end of the horizon which causes this policy to take much higher risk. This policy performs worse than the constant portfolio for that reason. Since there are no trading costs, the constant portfolio maintains the same level of position costlessly and manages the risk exposure better.

Figure 11 compares the optimal policy to the constant portfolio in the presence of trading costs. Both policies again have the same risk penalty by construction. Top-left panel illustrates that the difference in performance is more pronounced in the presence of trading costs. One reason for this is the excessive trading of the constant portfolio policy as illustrated in the medium-left and bottom panel. Compared to the previous case, we note that optimal policy trades much more slowly as shown in medium-right panel. The constant policy trades a lot after large return shocks in order to keep a constant dollar amount invested in the market portfolio. Therefore, the constant-dollar policy incurs much larger cumulative transaction costs than the optimal policy as we see in the bottom panel, which contributes a significant portion of the observed wealth difference between the two strategies.

Figure 12 compares the optimal policy to the buy-and-hold portfolio in the presence of trading costs. They both have the same risk penalty by construction. In the top-right panel, the starting position for the buy-and-hold policy is roughly $\$1.4 \times 10^9$. This policy performs better than the constant portfolio in this case as it never incurs trading costs. We note that the buy-hold policy is very slowly moving in building the position as it can never get out of the position to manage risk. This becomes the main driver of underperformance compared to the optimal policy.

Finally, Figure 13 compares the optimal policy to the myopic policy with transaction cost multiplier. Both policies have the same trading speeds but different aim portfolios. Since the risk penalties are not the same, wealth dynamics are not the main performance metric in this

case. For this reason, we also include the cumulative objective value which equals the difference between wealth and risk penalties. The performance difference as illustrated by objective values in the bottom-right panel is again substantial. The main driver seems to be excessive trading of the myopic policy. Since the myopic portfolio uses the conditional Markowitz portfolio as its aim position, it ends up trading a lot especially in the good state. Taking large positions, it also induces large risk penalties. This example shows the importance of accounting for the future dynamics of the state variables as this generates the difference between the aim portfolios of both policies.

5.6 Large vs. Small Portfolios

Managing transaction costs effectively will be very important when the portfolio size is large. In the absence of transaction costs, we know that the myopic portfolio, i.e., the conditional Markowitz portfolio, is optimal. Therefore, when the portfolio size is small, the difference between the optimal policy in the presence of transaction costs and the myopic portfolio may be very small. Since we are using realistic parameters, our model can also speak to the level of portfolio size at which managing transaction costs would provide significant benefits. For example, with $\gamma = 1 \times 10^{-10}$ we observe that our aim portfolios range from approximately \$20 billion to \$85 billion dollars.

Figure 14 compares the optimal policy to the myopic policy when $\gamma = 1 \times 10^{-5}$. In this case, the top-right panel tells us that the maximum aim portfolio across states is roughly \$2.8 million and in this case, there is no significant difference between performances.

Figure 15 compares the optimal policy to the myopic policy when $\gamma = 2.5 \times 10^{-8}$. With this calibration, the aim portfolios range from approximately \$20 million to \$900 million dollars. We observe that the myopic policy diverges a lot from the optimal policy by trading a lot and taking too much risk. It returns negative objective value and near-zero wealth levels. Thus, this simple exercise suggests that when the portfolio size is on the order of hundred millions, taking price impact into account is crucial.

5.7 Out-of-sample Analysis

The in-sample analysis was useful in studying the *expected* properties and benefits of a fully dynamic portfolio policy, but to better assess the value of the regime switching model, we perform an out-of-sample analysis. We implement a two-state regime switching model in this section for faster estimation of the parameters as we need to estimate a regime switching model every day from 1967 to 2017, roughly 12,600 estimations.

5.7.1 Calibration

First, we estimate the model parameters to determine the parameters of the objective function. We use all the available market return data from 1926 Q1 to 2017 Q2. Table 7 illustrates the estimated coefficients. We again observe that the expected return is lower in the high volatility state. The “good” state with higher expected return and low volatility is again more persistent.

Table 7: Parameter estimates for a two-state regime switching model using daily market return data from 1926 Q1 and 2017 Q2.

Parameter	Estimate	Parameter	Estimate
μ_1	0.0841%	σ_1	0.6110%
μ_2	-0.0955%	σ_2	1.8886%
P_{11}	0.9866	P_{12}	0.0134
P_{21}	0.0431	P_{22}	0.9569

Table 8: Transaction cost estimates in each regime labeled from the two-state regime switching model. λ_n denotes the transaction cost multiplier in regime n . The second column reports the results from a liquid subset in which we only include executions from stocks within the top 10% in market capitalization. Estimated values are multiplied by 10^{10} . Standard errors are double-clustered at the stock and calendar day level.

	<i>Dependent variable: IS</i>	
	<i>All Stocks</i>	<i>Liquid</i>
λ_1	1.772*** (0.255)	0.579*** (0.166)
λ_2	2.299*** (0.311)	1.254*** (0.335)

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

We estimate the transaction cost regimes using the same methodology, but now with two regimes. We again use the estimates from the liquid subset, i.e., the 50 stocks with largest market capitalizations. Formally, we run the following regression:

$$IS_i = \lambda_1 Q_i \mathbf{1}_{\{m(i)=1\}} + \lambda_2 Q_i \mathbf{1}_{\{m(i)=2\}} + \varepsilon_i$$

Table 8 illustrates the estimated coefficients. We observe that λ estimates are all highly significant. We find that λ_2 is greater than λ_1 and this difference is statically significant. Regime 2 has the lowest Sharpe ratio and thus can be interpreted as the distressed state.

5.7.2 Objective Function

The estimated two-state regime switching model and the calibrated transaction costs will determine the parameters of the out-of-sample objective function. Let x be any given policy. We will compute the out-of-sample performance of this policy by $W(x) - RP(x)$ where

$$W(x) = \sum_{t=1}^{T=12600} \rho^t \left[x_t r_{t+1} - \frac{1}{2} \lambda(s_t) (x_t - x_{t-1} R_t)^2 \right] \quad (26)$$

$$RP(x) = \sum_{t=1}^{T=12600} \frac{1}{2} \rho^t x_t^2 \sigma^2(s_t), \quad (27)$$

and s_t will equal to the state with the larger smoothed probability at time t , and σ and λ will be given by the calibrations in Table 7 and Table 8 (the liquid column), respectively.

The investor is not aware of the true parameters of the model and uses only information up to trading day t in order to make a trading decision for day $t + 1$, i.e., no policy will be able to use any forward looking data.

5.7.3 Optimal Policy

We construct our policy based on our theoretical analysis as follows. We will label this policy as the “optimal” policy as it is based on our dynamic model. First, we estimate a two-state regime switching model using the market return data from 1926 Q1 to 1967 Q2 (inclusive). We use these estimated parameters to construct a trading policy as formulated by Lemma 3. To apply our trading rule, we need to predict the regime of the next trading day. To accomplish this, we re-estimate a two-state regime switching model using return data from 1926 Q1 to the decision date. This estimation will provide smoothed probabilities for every trading day including the decision date. We will predict the next trading day’s regime using the state with the larger smoothed probability. For example, suppose that Regime 1’s smoothed probability for decision date is 0.52 and Regime 2’s smoothed probability for decision date is 0.48. We will predict the next trading day to be of Regime 1.

5.7.4 Benchmark Policies

We will use the constant portfolio and buy-and-hold portfolio as the benchmark policies.

We construct the constant portfolio policy in the out-of-sample data as follows. First, we estimate μ_{avg} and σ_{avg} using the market return data from 1926 Q1 to 1967 Q2. These parameters are held fixed throughout the investment horizon. The investor then constructs the following constant portfolio: $x_t^{\text{con}} = \frac{c\mu_{\text{avg}}}{\gamma\sigma_{\text{avg}}^2}$. We choose c so that the policy has the same risk exposure as the optimal policy.

The buy-and-hold portfolio is constructed similarly to its in-sample counterpart. The investor invests x_0 dollars (borrowed at the risk-free rate) into the market portfolio at the beginning of the investment horizon, i.e., on the first trading day of 1967 Q3, and then never trades but cumulates returns from its risky and risk-free asset positions. We choose x_0 so that the policy has the same total risk exposure as the optimal policy.

Figure 16 compares the optimal policy to the constant portfolio in the absence of trading costs in the out-of-sample data. The top-left panel illustrates that the optimal policy has higher performance in terms of terminal wealth. The results show that the regime-switching model captures predictability out-of sample and that it is valuable, absent transaction costs, to rebalance to time these regimes.

Figure 17 compares the optimal policy to the buy-and-hold portfolio in the absence of trading costs in the out-of-sample data. It confirms that the optimal policy outperforms the buy-and-hold portfolio out-of-sample in the absence of transaction costs.

Figure 18 compares the optimal policy to the constant portfolio in the presence of trading costs in the out-of-sample data. The top-left panel illustrates that the difference in performance is more pronounced in the presence of trading costs. The constant policy again trades a lot after large return shocks which reduces its overall performance. We can see that the difference in cumulative transaction costs paid by both strategies is very large and that this difference contributes substantially to the difference in wealth generated by both strategies. This hints to an interesting insight we confirm below. Even if expected return regimes are difficult to measure leading to a smaller out-of-sample performance in the absence of transaction costs, if transaction cost regimes are more accurately measured, which is plausible since t-costs vary with second moments, then optimally accounting for the variation in volatility and transaction costs leads to a sizable improvement in performance.

Figure 19 compares the optimal policy to the buy-and-hold portfolio in the presence of trading costs in the out-of-sample data. We find that the outperformance of the optimal policy is again substantial. Here the myopic policy again builds the position very slowly but ends-up with a very large a position at the end of the sample which increases the total risk. In the top-right panel, the starting position for the buy-and-hold policy is roughly $\$4.7 \times 10^9$. This is substantially lower than the aim portfolio of the optimal policy in the low-volatility state.

Overall, this out-of-sample analysis illustrates that the outperformance of the optimal policy is robust to parameter uncertainty of the regime switching model.

5.8 Which parameter should you time?

In this section, we investigate the value of timing each switching parameter of the general model. The switching parameters are μ , σ and λ . It is well-known, at least since Merton (1980), that expected returns are estimated less precisely than volatilities. Further, Moreira and Muir (2017) have shown that there are gains to scaling down the risky asset exposure in response to an increase in the market's variance, which suggests that the conditional mean of the market moves less than one-for-one with its variance. One might thus expect that out-of-sample the benefits of timing changes in volatility could be larger than timing changes in expected returns. We show some evidence to that effect below. Further, since transaction costs vary with volatilities, we also provide quantitative evidence about the value of timing transaction cost regimes.

We use the the implementation of the optimal policy from the out-of-sample analysis to account for the potential bias introduced by imprecisely estimated parameters. First, we study the value of timing the switches in either volatility or expected returns in the absence of trading costs. In this analysis, if the investor times volatility, he takes into account that the volatility is time-varying between two states but assumes that expected return is constant throughout the investment horizon and is given by μ_{avg} (as in the case of the constant portfolio rule). Similarly, if the investor times expected returns, he models them as time-varying between high and low states and internalizes the potential switches in the expected return in his trading rule but he assumes that the volatility stays constant at a level of σ_{avg} (as in the case of the constant portfolio rule). We scale the policies so that they take the same risk.

Figure 20 compares these two timing approaches in the absence of trading costs using an out-of-sample trading approach. We scale both policies so that they both have the same risk exposure as the optimal policy that times both parameters. We find that timing volatility provides much higher performance. The terminal wealth of the policy that only times volatility is actually higher than the terminal wealth of the optimal policy that times both parameters as shown in Figure 16. This illustrates that trying to time expected returns may be actually detrimental in an out-of-sample trading strategy. The top-right panel shows that the μ -timing policy has a wider range of positions compared to the range observed in the σ -timing policy. In the absence of t-costs the strategies switch to their conditional mean-variance Markowitz portfolios in every state. Recall that the estimated mean in the state 2 is negative and the volatility is high. This implies that the μ -timing strategy, which underestimates the volatility in that regime, takes a very large short position in the risky asset. This hurts the out-of-sample performance of the strategy relative to the volatility timing strategy, probably because the negative expected return in those states is not precisely estimated.

If there are trading costs in the model, then λ will be switching through time between high and low transaction cost regimes. If an investor does not time the switches in λ , then the investor uses an unconditional average of λ_{avg} which is estimated from running following regression in the liquid subset:

$$IS_i = \lambda_{\text{avg}} Q_i + \varepsilon_i$$

where Q_i is the dollar size of the order. We estimated λ_{avg} to be 0.766×10^{-10} which is between λ_1 and λ_2 , as expected.

Now we consider combined timing strategies: Timing σ and μ , timing σ and λ or timing μ and λ . In all three timing strategies, the left-out parameter is set to its unconditional average. We consider the comparison across these policies in two different assumptions of γ : high risk-aversion and low risk-aversion.²⁰ Figure 21 compares these three policies in the presence of transaction costs in the high risk-aversion case. We observe that the top performing policy times σ and λ and the worst performing policy times σ and μ .

Figure 22 compares these three policies in the low risk-aversion case. We again observe that the worst performing policy times σ and μ but the underperformance is economically smaller. This underscores that the benefits from timing volatility and transaction costs become more important when the size of the portfolio is large.

6 Conclusion

In this paper, we develop a closed-form solution for the dynamic asset allocation when expected returns, covariances, and price impact parameters follow a multi-state regime switching model. Under mean-variance objective function, we compute the optimal trading rule of the investor analytically by characterizing the trading speed and aim portfolio. Specifically, the aim portfolio is a weighted average of the conditional Markowitz portfolios in all potential states. The weight on each conditional Markowitz portfolio depends on the likelihood of transitioning to that state, the state's persistence, the risk, and transaction costs faced in that state compared to the current one. Similarly, the optimal trading speed is a function of the relative magnitude of the transaction costs in various states and their transition probabilities. One of the significant implications of our model is that the optimal portfolio can deviate substantially from the conditional Markowitz portfolio in anticipation of possible future shifts in relative risk and/or transaction costs.

We show that the model is equally tractable when price changes or returns follow a regime-switching model. The latter aligns better with the empirical dynamics of asset returns. We utilize this framework to optimally time the broad value-weighted market portfolio, accounting for time-varying expected returns, volatility, and transaction costs. We use a large proprietary data on institutional trading costs to estimate the price impact parameters. We find that trading costs vary significantly across regimes and tend to be higher as market volatility increases.

We test our trading strategy both in-sample and out-of-sample and find that there are substantial benefits to the use of our approach. For the out-of-sample test, the state probabilities are estimated using only data in the information set of an agent on the day preceding the trading date. We compare the performance of our optimal dynamic strategy to various benchmarks: a constant dollar investment in the risky asset, a buy-and-hold portfolio, and a myopic policy with

²⁰Note that in the absence of trading costs, changing risk aversion would not matter, as the wealth values will just be scaled by the ratio of the risk-aversion parameters.

optimal trading speeds borrowed from the optimal solution. Our dynamic strategy outperforms all of these alternatives significantly. Out-of-sample, the benefits of timing volatility and transaction costs dominate those of timing expected returns, especially when assets under management are sizable.

References

- Admati, Anat, and Paul Pfleiderer, 1988, A theory of intraday patterns: Volume and price variability, *Review of Financial Studies* 1, 3–40.
- Ang, Andrew, and Geert Bekaert, 2002, International asset allocation with regime shifts, *The Review of Financial Studies* 15, 1137–1187.
- Ang, Andrew, and Allan Timmermann, 2012, Regime changes and financial markets, *Annu. Rev. Financ. Econ.* 4, 313–337.
- Asness, Clifford, Andrea Frazzini, and Lasse Pedersen, 2012, Leverage aversion and risk parity, *Financial Analysts Journal* 68, 47–59.
- Balduzzi, Pierluigi, and Anthony W Lynch, 1999, Transaction costs and predictability: Some utility cost calculations, *Journal of Financial Economics* 52, 47–78.
- Brennan, Michael J., Eduardo S. Schwartz, and Ronald Lagnado, 1997, Strategic asset allocation, *Journal of Economic Dynamics and Control* 21, 1377–1403.
- Bridgewater, 2011, Risk parity is about balance, *Bridgewater Associates Research Note*.
- 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 Y, Yeung Lewis Chan, and Luis M Viceira, 2003, A multivariate model of strategic asset allocation, *Journal of Financial Economics* 67, 41–80.
- Campbell, John Y, and Luis M Viceira, 2002, *Strategic asset allocation: portfolio choice for long-term investors*. (Oxford University Press Oxford).
- Cochrane, John H, 2007, Portfolio theory, University of Chicago Working Paper.
- Constantinides, George M., 1986, Capital market equilibrium with transaction costs, *The Journal of Political Economy* pp. 842–862.
- Cvitanić, Jakša, 2001, Theory of portfolio optimization in markets with frictions, in Elyès Jouini, Jakša Cvitanic, and Marek Musiela, eds.: *Option Pricing, Interest Rates and Risk Management* (Cambridge Univ Press, Cambridge, UK).
- Davis, M. H. A., and A. R. Norman, 1990, Portfolio selection with transaction costs, *Mathematics of Operations Research* 15, 676–713.
- Detemple, Jérôme, and Marcel Rindisbacher, 2010, Dynamic asset allocation: Portfolio decomposition formula and applications, *Review of Financial Studies* 23, 25–100.

- 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.
- Foster, F Douglas, and Subramanian Viswanathan, 1993, Variations in trading volume, return volatility, and trading costs; evidence on recent price formation models, *The Journal of Finance* 48, 187–211.
- Gârleanu, Nicolae, and Lasse H. Pedersen, 2013, Dynamic trading with predictable returns and transaction costs, *The Journal of Finance* 68, 2309–2340.
- Grinold, Richard C., and Ronald N. Kahn, 1999, *Active portfolio management: a quantitative approach for providing superior returns and controlling risk*. (McGraw-Hill Companies New York, NY).
- Guidolin, Massimo, and Allan Timmermann, 2006, An econometric model of nonlinear dynamics in the joint distribution of stock and bond returns, *Journal of Applied Econometrics* 21, 1–22.
- Hamilton, James D, and Raul Susmel, 1994, Autoregressive conditional heteroskedasticity and changes in regime, *Journal of econometrics* 64, 307–333.
- Hugonnier, Julien, Florian Pelgrin, and Pascal St-Amour, 2012, Health and (other) asset holdings, *The Review of Economic Studies* 0, 1–48.
- Jang, Bong-Gyu, Hyeng Keun Koo, Hong Liu, and Mark Loewenstein, 2007, Liquidity premia and transaction costs, *The Journal of Finance* 62, 2329–2366.
- Kim, Tong Suk, and Edward Omberg, 1996, Dynamic nonmyopic portfolio behavior, *Review of Financial Studies* 9, 141–161.
- Litterman, Robert, 2005, Multi-period portfolio optimization, working paper.
- Liu, Jun, 2007, Portfolio selection in stochastic environments, *Review of Financial Studies* 20, 1–39.
- Ljungqvist, Lars, and Thomas J. Sargent, 2004, *Recursive Macroeconomic Theory*. (The MIT press Cambridge, MA).
- 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.
- Markowitz, Harry M., 1952, Portfolio selection, *Journal of Finance* 7, 77–91.
- Mei, Xiaoling, Victor DeMiguel, and Francisco J Nogales, 2016, Multiperiod portfolio optimization with multiple risky assets and general transaction costs, *Journal of Banking & Finance* 69, 108–120.
- 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.
- Merton, Robert C., 1980, On estimating the expected return on the market, *Journal of Financial Economics* 8, 323–361.
- Moreira, Alan, and Tyler Muir, 2017, Volatility-managed portfolios, *The Journal of Finance* 72, 1611–1644.
- Perold, Andre F., 1988, The implementation shortfall: Paper versus reality, *The Journal of Portfolio Management* 14, 4–9.
- Rosenberg, Barr, 1972, The behavior of random variables with nonstationary variance and the distribution of stock prices, Working paper University of California Berkeley.
- 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, eds.: *Handbooks in Operations Research and Management Science* (Elsevier B.V., Amsterdam).
- Stoll, Hans R., 1978, The supply of dealer services in securities markets, *The Journal of Finance* 33, 1133–1151.

Figure 6: Regimes. The first four plots illustrate the corresponding smoothed probabilities for each regime. We use the following color codes: Green represents regime 1, blue represents regime 2, yellow represents regime 3, and red represents regime 4. The fifth plot illustrates the color-coded regimes vertically by using the maximum smoothed probability for identification. Market returns are in black.

Figure 7: Unconditional Markowitz portfolio and aim portfolios in the absence and presence of trading costs. In the top panel, trading costs are assumed to be zero, thus the aim portfolio is equal to the conditional Markowitz portfolio. We also plot the unconditional Markowitz portfolio which we label as the “Constant” portfolio. In the bottom panel, trading costs are set according to Table 5 using executions from the liquid subset. Coefficient of risk aversion is given by $\gamma = 1 \times 10^{-10}$ (can be thought as corresponding to a relative risk aversion of 1 for an agent with \$10 billion dollars under management.).

Figure 8: Trading speed across different regimes. Trading costs are set according to Table 5 using executions on very large-cap stocks. Coefficient of risk aversion is given by $\gamma = 1 \times 10^{-10}$ (can be thought as corresponding to a relative risk aversion of 1 for an agent with \$10 billion dollars under management.).

Figure 9: This figure compares the in-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). 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.

Figure 10: This figure compares the in-sample performance of the optimal policy with a buy-and-hold 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). 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.

Figure 11: This figure compares the in-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). 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.

Figure 12: This figure compares the in-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). 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.

Figure 13: This figure compares the in-sample performance of the optimal policy with a myopic policy in the presence of trading costs. Both strategies start from zero-wealth. We adjust the myopic policy so that it has the same trading speed as the optimal policy. Top-left panel shows the cumulative wealth of each policy. 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-left panel illustrates the cumulative cost of these trades. Bottom-right panel shows the realized objective value of each strategy, which is the main comparison metric, by subtracting the risk penalty from the wealth generated.

Figure 14: This figure compares the in-sample performance of the optimal policy with a conditional Markowitz portfolio for a small investor ($\gamma = 1 \times 10^{-5}$) in the presence of trading costs. Recall that conditional Markowitz portfolio is optimal in the absence of trading costs. Both strategies start from zero-wealth. Top-left panel shows the cumulative wealth of each policy. 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-left panel illustrates the cumulative cost of these trades. Bottom-right panel shows the realized objective value of each strategy, which is the main comparison metric, by subtracting the risk penalty from the wealth generated.

Figure 15: This figure compares the in-sample performance of the optimal policy with a conditional Markowitz portfolio for a medium-size investor ($\gamma = 2.5 \times 10^{-8}$) in the presence of trading costs. Recall that conditional Markowitz portfolio is optimal in the absence of trading costs. Both strategies start from zero-wealth. Top-left panel shows the cumulative wealth of each policy. 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-left panel illustrates the cumulative cost of these trades. Bottom-right panel shows the realized objective value of each strategy, which is the main comparison metric, by subtracting the risk penalty from the wealth generated.

Figure 16: 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.

Figure 17: 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.

Figure 18: 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.

Figure 19: 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.

Figure 20: 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.

Figure 21: 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} .

Figure 22: 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} .

A Appendix

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 and C, respectively.

In Appendix D 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 E, 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 F we construct these preference specifications in a recursive framework, and show in Appendix G, 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.

B 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:

$$dS_t = \mu(s_t)dt + \sigma(s_t)dZ_t \quad (28)$$

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.²¹

²¹In 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

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

$$\max_{\theta_t} \mathbb{E} \left[\int_{t=0}^{\infty} \left\{ n_t^\top \mu(s_t) - \frac{1}{2} \gamma n_t^\top \Sigma(s_t) n_t - \frac{1}{2} \theta_t^\top \Lambda(s_t) \theta_t \right\} e^{-\rho t} dt \right] \quad (29)$$

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.²²

We interpret this as an investor who maximizes his expected wealth, $\mathbb{E}[W_\tau]$ 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

$$dW_t = n_t dS_t - \frac{1}{2} \theta_t^\top \Lambda(s_t) \theta_t dt - \frac{1}{2} \gamma n_t^\top \Sigma(s_t) n_t dt, \quad (30)$$

$$dn_t = \theta_t dt. \quad (31)$$

Note that $\mathbb{E}[W_\tau] = \int_0^\infty \rho e^{-\rho t} W_t dt = \int_0^\infty e^{-\rho t} dW_t$ if the transversality condition $\lim_{T \rightarrow \infty} \mathbb{E}[e^{-\rho T} W_T] = 0$ holds. We solve this problem using dynamic programming. Define the value function:

$$J^s(n_t) = \max_{\theta_u} \mathbb{E} \left[\int_{u=t}^{\infty} e^{-\rho(u-t)} \left\{ n_u^\top \mu(s_u) - \frac{1}{2} \gamma n_u^\top \Sigma(s_u) n_u - \frac{1}{2} \theta_u^\top \Lambda(s_u) \theta_u \right\} du \mid s_t = s \right]. \quad (32)$$

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

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

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.

²²Naturally 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.

Plugging this back into the HJB equation we get:

$$0 = n^\top \mu_s - \frac{1}{2} \gamma n^\top \Sigma_s n + \frac{1}{2} (\nabla J^s)^\top \Lambda_s^{-1} \nabla J^s + \pi_{s,s'} (J^{s'} - J^s) - \rho J^s \quad (33)$$

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.

C Continuous time model of price changes: Finite horizon

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

$$\max_{\theta_t} \mathbb{E} \left[\int_{t=0}^T \left\{ n_t^\top \mu(s_t) - \frac{1}{2} \gamma n_t^\top \Sigma(s_t) n_t - \frac{1}{2} \theta_t^\top \Lambda(s_t) \theta_t \right\} dt \right] \quad (34)$$

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, $\mathbb{E}[W_T]$ evaluated at a fixed time T and faces quadratic transaction costs and holding costs, so that her wealth dynamics are

given by:

$$dW_t = n_t dS_t - \frac{1}{2} \theta_t^\top \Lambda(s_t) \theta_t dt - \frac{1}{2} \gamma n_t^\top \Sigma(s_t) n_t dt \quad (35)$$

$$dn_t = \theta_t dt \quad (36)$$

Define the value function:

$$J^s(n_t, t) = \max_{\theta_u} \mathbb{E} \left[\int_{u=t}^T \left\{ n_u^\top \mu(s_u) - \frac{1}{2} \gamma n_u^\top \Sigma(s_u) n_u - \frac{1}{2} \theta_u^\top \Lambda(s_u) \theta_u \right\} du \mid s_t = s \right] \quad (37)$$

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 time derivative.

The first order condition is

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

Plugging back into the HJB equation we get:

$$0 = n^\top \mu_s - \frac{1}{2} \gamma n^\top \Sigma_s n + \frac{1}{2} (\nabla J^s)^\top \Lambda_s^{-1} \nabla J^s + \pi_{s,s'}(J^{s'} - J^s) + \dot{J}^s \quad (38)$$

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

$$\begin{aligned} 0 = & -\frac{1}{2} n^\top \{ Q_s \Lambda_s^{-1} Q_s + \gamma \Sigma_s - \dot{Q}_s + \pi_{s,s'}(Q_{s'} - Q_s) \} n \\ & + n^\top \{ \mu_s - Q_s \Lambda_s^{-1} q_s - \dot{q}_s + \pi_{s,s'}(q_{s'} - q_s) \} + \frac{1}{2} q_s^\top \Lambda_s^{-1} q_s - \dot{c}_s + \pi_{s,s'}(c_{s'} - c_s) \end{aligned} \quad (39)$$

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

$$\dot{Q}_s = Q_s \Lambda_s^{-1} Q_s + \gamma \Sigma_s - Q_s + \pi_{s,s'}(Q_{s'} - Q_s) \quad (40)$$

$$\dot{q}_s = \mu_s - Q_s \Lambda_s^{-1} q_s - q_s + \pi_{s,s'}(q_{s'} - q_s) \quad (41)$$

$$\dot{c}_s = \frac{1}{2} q_s^\top \Lambda_s^{-1} q_s - c_s + \pi_{s,s'}(c_{s'} - c_s) \quad (42)$$

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.

D 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:

$$dW_t = n_t dS_t - \frac{1}{2} \theta_t^\top \Lambda(s_t) \theta_t dt \quad (43)$$

$$dn_t = \theta_t dt \quad (44)$$

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

$$\max_{\theta_t} \mathbb{E} [-e^{-\gamma W_T}] \quad (45)$$

Define the ‘value’ function:

$$J^s(n_t, t) = \max_{\theta_u} \mathbb{E} [-e^{-\gamma \int_t^T dW_u} \mid s_t = s] \quad (46)$$

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 (38) 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 (40)-(42) 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).

E 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 $M_t = E_t[U(W_T)]$ 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

$$dH_t = \mu_H dt + \sigma_H dZ_t + \eta_H (dN_t - \lambda_t dt) \quad (47)$$

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:

$$\mu_H = -\frac{1}{2} \frac{U''(H)}{U'(H)} \sigma_H^2 + \lambda(\eta_H - \frac{U(H + \eta_H) - U(H)}{U'(H)}) \quad (48)$$

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:

$$H_t = E_t[W_T - \int_t^T \mu_H(H_t, \sigma_H, \eta_H) dt] \quad (49)$$

where the driver of the BSDE is given in equation (48) 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 $E[U(W_T)]$, 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 (48-49) 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:

$$H_t = E_t[W_T - \int_t^T \mu_H(H_t, \sigma_H, \eta_H) dt] \quad (50)$$

$$\mu_H = -\frac{1}{2} \frac{U_1''(H)}{U_1'(H)} \sigma_H^2 + \lambda \left(\eta_H - \frac{U_2(H + \eta_H) - U_2(H)}{U_2'(H)} \right) \quad (51)$$

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).²³

F 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 (47). At any time $t < T$ the certainty equivalent is defined by the following recursion

$$W(H_t, 0, 0) = E_t[W(H_t + \mu_H dt, \sigma_H dZ_t, \eta_H(dN_t - \lambda)dt)] \quad (52)$$

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:

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

Using this we can rewrite the recursion (52), 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 (50) above.

²³The only difference is that we do not have any intermediate consumption, and do not restrict to CRRA utility functions.

G Source-Dependent SDU with Vanishing Jump Risk

We will now show that our objective function in Appendix C (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 $\gamma_2 \rightarrow 0$. Indeed, from the definition in (50) the certainty equivalent of such an agent is the solution (H, σ_H, η_H) of the BSDE:

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

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

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

Further, with W_t dynamics given in (43) 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 $\sigma_H = \sigma_W$):

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

This is the objective function we considered in Appendix C.

G.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.