Equivalence of Learning Technologies Between Rotated and Unrotated Versions of the Model

Marco Sammon

February 12, 2021

1 Introduction

In the baseline version of the model, assets have correlated payoffs because of their common exposure to the systematic risk-factor. In addition, investors receive correlated signals about the payoffs of the *assets*, rather than the payoffs of the underlying orthogonal *risk-factors*. When the ETF is present, however, the number of risk-factors is equal to the number of assets. This condition implies that there exists an equivalent economy where asset payoffs and signals are orthogonal (see e.g., Appendix B of Kacperczyk, Van Nieuwerburgh and Veldkamp [2016]).

The existence of this equivalent economy means the assumption of correlated signals has no effect on investors' optimal attention allocation. The intuition for this claim is that any investor could rotate the set of correlated signals/payoffs to a set of orthogonal signals/payoffs, and back out an independent signal about each risk-factor.

When the ETF is not present, the learning technology and asset payoffs are written as though there are *more* risk-factors than assets. There is, however, an equivalent economy where the number of risk-factors is *equal* to the number of assets. This can be accomplished by e.g., (1) removing the systematic risk-factor f from each stock's payoff, and (2) rather than having the stock-specific risk-factors η_i be independent, make them correlated such that $cov(\eta_i, \eta_j) = \sigma_f^2$ where σ_f^2 is the volatility of the systematic risk-factor.

When the ETF is not present, the assumption of correlated signals may affect the optimal attention allocation. Although investors could rotate the economy to a set of orthogonal signals/payoffs, there is no way to isolate an independent signal about what I label as the systematic risk-factor f i.e., the common component in the η_i's in the equivalent economy with an equal number of risk-factors and assets.

Despite this, I solve the model numerically with correlated assets and signals, instead of rotating the economy and using the closed-form solutions in Kaperczyk et. al. [2016]. This is to ensure that the total attention allocated by investors is equal between economies with and without the ETF[1]. This note explains why my solution method preserves total attention, and shows examples where the rotation method may not.

This note is organized into three parts: Section 2 outlines the general mapping between my numerical solution method, and the closed-form solution in Kacperczyk et. al. [2016]. Section 3 outlines the exact steps I use to map the rotated economy (i.e. the economy with orthogonal assets/signals) to the un-rotated economy (i.e. the economy with correlated assets/signals). Section 4 walks through two examples where the learning technologies are not equivalent between the rotated and unrotated versions of the economy.

2 General Mapping

Even when the ETF is present, investors get signals about the payoffs of the underlying *assets* rather than the payoffs of the underlying *risk-factors*. The

[1] Another reason this assumption is tractability: According to Admati (1985), there is no closed-form solution for prices in the scenario where investors receive an independent signal about the systematic risk-factor, but cannot trade on it directly i.e. there are more risk-factors with *independent* signals than assets.

attention allocation problem is solved numerically assuming investors receive these correlated signals. The model can also be solved by: (1) rotating the model to have orthogonal asset payoffs and signals (2) using the formulas in Kacperczyk et. al. [2016] to find the optimal attention allocation and (3) rotating the economy back to the original covariance strcture.

Numerical methods, however, would still be required to find the rotation matrix and find the corresponding total attention constraint between the rotated and unrotated versions of the economy. Here are the steps in that procedure:

1. Choose some range for the total attention constraint in the rotated version of the model. Loop over every $\hat{K}$ between some lower-bound $\underline{K}$ and some upper-bound $\overline{K}$.
2. For each of these $\hat{K}$'s, determine what the optimal attention allocation would be if each stock had a β_i of zero on the systematic risk-factor f.
3. Using an eigendecomposition, find the rotation matrix Q that maps the economy with orthogonal asset payoffs and signals to the economy with correlated asset payoffs and signals.
4. Choose the $\hat{K}$ where the total attention constraint is satisfied in the unrotated version of the economy. In the case of $K = 1$, choose $\hat{K}$ such that after applying the rotation factor Q, the attention allocations, i.e. the K_i's, add up to one.

3 Specific Procedure

Here is a more detailed version of the procedure outlined above:

1. Choose some range for the total attention constraint in the rotated version of the model, $\hat{K}$, say between 0.025 and 2, in increments of 0.025 for $K = 1$.

2. For each of these $\hat{K}$'s, determine what the optimal attention allocation would be if all the assets were independent i.e., if all the stocks had a β_i of zero on the systematic risk-factor f. This is given by the formulas in Kacperczyk et. al. [2016]. Call these optimal attention allocations K_i^*.

3. For any (not necessarily optimal) attention allocation, define L as a diagonal matrix with $1/K_i$ in every (i,i) entry. Define L^* as a particular case of L when using the K_i^*'s. With independent risk-factors and signals, $L^* = \Sigma_e$, where Σ_e denotes the variance-covariance matrix of signal noises.

4. Risk-factors and signals are not independent, so L^* is not necessarily equal to Σ_e. Instead, Σ_e is equal to $\Gamma L \Gamma$, where Γ is a matrix defining each asset's exposure to the systematic risk-factor i.e. an identity matrix with an extra column containing the β_i's. All the β_i's in the baseline version of the model are assumed to be one. L here can be any (possibly non-optimal) attention allocation in the unrotated version of the economy.

5. For every L feasible with the total attention constraint K (in the baseline specification $K = 1$), do an eigendecomposition of $\hat{\Sigma_e} = \Gamma L \Gamma'$ into $G \Lambda G'$ (note that Γ is not usually equal to G because L's diagonal elements are not usually the eigenvalues of $\hat{\Sigma_e}$) to solve for the rotation factor G. Define a function which returns the normalized difference between GL^*G' and $\hat{\Sigma_e}$. This value, which I call diff(L), will be equal to zero if L in the unrotated version of the economy maps to the optimal attention allocation L^* in the rotated version of the economy.

   - Note: L^* is not necessarily equal to L because L^* is from the orthogonal version of the economy, while Σ_e is from the non-orthogonal version of the economy.

6. For each total attention allocation looped over, $\hat{K}$, find the K_i's that minimize diff(L).
   - This is a non-linear problem, so try many starting points when doing this optimization to avoid getting stuck at a local minimum.
7. Find the $\hat{K}$ to minimize the distance between the sum of the K_i's in L and the total information constraint K (which is set to 1 in the baseline). If this distance is zero, there is an equivalence between the learning capacity K in the unrotated economy and $\hat{K}$ in the rotated economy.
   - Whether $\hat{K}$ is bigger or smaller than K depends on the risk-bearing capacity of the economy. If there is a lot of risk bearing capacity, $\hat{K}$ will tend to be bigger than K. Otherwise, $\hat{K}$ will tend to be smaller than K.

4 Numerical Examples

There is no guarantee that the sum of the $K_i's$ in L^* (the optimal attention allocation in the rotated economy) are the same as the sum of the $K_i's$ in L (the optimal attention allocation in the unrotated economy). This makes it difficult to compare total learning capacities (K's) between rotated and unrotated economies. Here are two numerical examples of this phenomenon:

1. Suppose the share of informed investors is 50%, investor risk aversion $\rho = 0.05$, the volatility of the systematic risk factor $\sigma_f = 0.05$ and total attention $K = 1$. The corresponding $\hat{K}$ in the rotated economy is 1.125, so in the rotated economy, we need to give investors more total attention to allocate if we want things to be equivalent to the unrotated economy.

2. Suppose the share of informed investors is 50%, investor risk aversion $\rho = 0.15$, the volatility of the systematic risk factor $\sigma_f = 0.25$ and total learning capacity $K = 1$. The $\hat{K}$ in the rotated economy is 0.175, so in the rotated economy, we need to give investors less total attention to allocate if we want things to be equivalent to the unrotated economy.

In both these cases, the total attention capacity is different in the unrotated economy and the rotated economy. This illustrates why it is not meaningful to compare total attention capacities across different rotated economies. Solving the unrotated model is one way to ensure that the total attention capacity is equal across the rotated and unrotated economies. As a result, solving the unrotated model sidesteps the fact that rotation factors (and therefore equivalent total learning capacities) will be different for economies with and without the ETF, even though all the other parameters are equal.