Midlincoln Research

February - 6 2025

Analyst: Ovanes Oganisian

A crucial distinction: forward vs inverse problems

The forward problem (well-defined)

The inverse problem (the subject of this note)

Why regression can succeed — and why that does not solve the problem

Ranking as a projective map and equivalence classes

Why multiple weight vectors produce the same ranking

Zero-projection stocks and the decision hyperplane

Why repeated years help — but do not restore uniqueness

Monte Carlo as geometry exploration, not optimization

A concrete compute perspective
Why this behavior was historically misunderstood

What this reframes about factor weights

A precise statement of the core insight

Quants

Why Learning Factor Weights Is an Ill-Posed Inverse Problem

Ranking, equivalence classes, and why robustness beats optimization

In the first note (Ranking Before Prediction), I argued that many quantitative equity systems are not prediction engines but **ranking systems**. They normalize information, combine factors, induce an ordering, and act almost exclusively on the extremes of that ordering.

This naturally leads to the next question:

> If ranking is the object, why not *learn* the factor weights that produce the correct ranking?

At first glance, this seems straightforward.
Given factor information at time t and realized returns over a future window, why not infer the weights that would have generated the observed ordering?

This note explains why that question — **weight recovery from realized rankings** — is not well-behaved in the usual sense, why this difficulty is structural rather than technical, and why practitioners repeatedly converge on heuristics, constraints, Monte Carlo exploration, and stability analysis instead of precise coefficient estimation.

A crucial distinction: forward vs inverse problems

Let's start by separating two problems that are often conflated.

The forward problem (well-defined)

Given:

- a set of factor vectors Z_i,
- a weight vector a,

we compute scores:

$$s_i = a^\top Z_i$$

and then rank the stocks by s_i.

This is:

- deterministic,
- stable,
- and unambiguous.

There is nothing ill-defined about producing a ranking from weights.

The inverse problem (the subject of this note)

Given:

- factor vectors Z_i,
- a realized return ordering over some future window,

infer the weights a that "generated" that ordering.

This is a different problem entirely.

It is this **inverse problem** — mapping from an observed permutation back to weights — that behaves badly.

Why regression can succeed — and why that does not solve the problem

Linear regression is often the first tool considered.

Regression minimizes average squared error:

$$\min_a \sum_i (r_i - a^\top Z_i)^2$$

Regression **can** recover a ranking that matches realized returns when:

- signal is strong and linear,
- noise is low,
- horizons are long,
- the universe is small or homogeneous.

In those regimes, many methods agree.

But regression is not *trying* to recover an ordering.
It optimizes level fit, not rank stability, truncation, or tail behavior.

When regression succeeds at ranking, it does so incidentally — not structurally.

That distinction matters because the regimes that motivate ranking systems (short horizons, heavy tails, large universes) are exactly those where regression's success probability collapses.

Ranking as a projective map and equivalence classes

The key mathematical structure is simple but easy to misstate.

If a weight vector a produces a ranking, then:

$$\text{any } c\cdot a \text{ for any} c > 0$$

produces the **same ranking**.

This does **not** mean:

- signs alone matter,
- magnitudes are irrelevant,
- all weights are equivalent.

It means only this:

> **Ranking is invariant to global rescaling.**

Global scale removes one degree of freedom.
What remains — the **relative proportions between factors** — is essential.

Vectors like (0.1,0.7) and (0.3,0.4) have the same signs but different directions.
They lie on different rays and can produce very different rankings in practice — exactly as practitioners observe.

So the natural object is not a point in weight space, but a **direction**: a ray modulo scale.

Why multiple weight vectors produce the same ranking

Once we work in direction space, a deeper issue appears.

Each pairwise ordering constraint:

$$a^{\top}(Z_i - Z_{j)} > 0$$

defines a half-space in weight space.

The set of weights consistent with an observed ranking is therefore:

- an intersection of half-spaces,
- a **region**, not a point.

This region reflects:

- correlated factors,
- near-ties in returns,
- noise in outcomes,
- truncation of relevance to the top of the ranking.

There is no mathematical reason for this region to collapse to a unique direction — and empirically, it rarely does.

Zero-projection stocks and the decision hyperplane

An intuition that often arises in practice is the importance of stocks with **near-zero projected returns**.

Geometrically, for a given weight direction a, the set:

$$a^{\top}Z=0$$

defines a **decision hyperplane**.

Stocks near this hyperplane:

- flip sign under small perturbations,
- reshuffle easily in the middle of the ranking,
- dominate turnover and instability.

Permutations among these stocks do not materially affect the coarse ranking.

This is not a kernel in the strict algebraic sense — the ranking operator is not linear — but it is a meaningful **equivalence structure**: a region of indifference where ordering carries little economic meaning.

Studying this boundary often reveals more about robustness than studying precise coefficient values.

Why repeated years help — but do not restore uniqueness

A natural response to non-uniqueness is:

> “Test the weights across multiple years.”

This is sensible and necessary — but it does something specific.

Multiple periods:

- intersect feasible regions,
- eliminate fragile directions,
- stabilize signs and ratios.

What they typically do **not** do:

- identify a unique optimal weight vector,
- eliminate trade-off directions,
- collapse the solution to a point.

What survives is usually a **cone or tube of directions** — narrower, more stable, but still not unique.

This is not failure.
It is the correct outcome given the object.

Monte Carlo as geometry exploration, not optimization

Seen this way, Monte Carlo sampling is not a crude optimization trick.

Properly used, it:

- samples directions (after removing scale),
- identifies which directions are compatible with observed rankings,
- reveals sign stability and factor trade-offs,
- exposes the width of the feasible region.

Monte Carlo answers:

> "What must be true for the ordering to hold?"

not:

> "What is the one true set of weights?"

That shift in question is essential.

A concrete compute perspective

To make this tangible, consider a realistic setup:

- **3000 stocks**
- **60 countries**
- **10 sectors**
- ~5–10 normalized factors per stock

A typical workflow might be:

1. Normalize factors within sector.
2. Rank stocks within each country by a linear score.
3. Select top and bottom subsets.

Now suppose we explore weight directions via Monte Carlo.

For each sampled direction:

- compute dot products (3000×10),

- sort stocks within each country,
- evaluate a ranking metric (e.g. top-k overlap or average top-decile return).

Even with tens of thousands of samples:

- the dot products dominate compute,
- sorting is secondary,
- the problem is tractable on ordinary hardware.

But note what compute is **not** doing here:

- it is not converging to a unique solution,
- it is not “learning” coefficients in the regression sense.

Compute is being used to **map geometry**, not to chase precision.

Why this behavior was historically misunderstood

In the early 2000s, treating realized rankings as mathematical objects felt unnatural.

Statistics focused on expectations.
Algebra focused on abstract permutations, not empirical ones.
Finance focused on continuous-time models.

Ranking under uncertainty had no natural home.

Only later — driven by information retrieval, recommendation systems, and machine learning — did ranking emerge as a first-class object of study.

Practitioners encountered it earlier because they had no choice: capital allocation is discrete, truncated, and tail-driven.

Theory followed practice, not the other way around.

What this reframes about factor weights

The conclusion is not that factor models are meaningless.

It is more subtle:

- factor weights are **coordinates**, not parameters of reality;
- their global scale is irrelevant;
- their relative proportions matter;
- their exact values are rarely identifiable from rankings alone.

The real objects of interest are:

- stable directions,
- sign constraints,
- trade-off structure,
- robustness under perturbation.

This explains why experienced practitioners:

- stop short of aggressive optimization,
- rely on heuristics and constraints,
- emphasize turnover and stability,
- distrust “best” models.

A precise statement of the core insight

Here is the cleanest formulation:

> While inducing a ranking from factor weights is a well-defined forward problem, inferring factor weights from a realized return ordering is a non-identifiable inverse problem. The ranking map identifies equivalence classes of weight directions, and repeated observations typically stabilize regions rather than single solutions.

That statement avoids drama and matches both mathematics and practice.

Closing thought

The persistent failure to pin down “true” factor weights is not a tooling problem.

It is a signal that the object being studied — ranking under noise and truncation — does not support that level of identification.

Once this is accepted, the focus shifts naturally:
from optimization to robustness,
from precision to geometry,
from point estimates to regions.

That shift is not a concession.
It is progress.

The information contained in this document is neither an offer to sell nor a solicitation of an offer to purchase interests in any referenced investment nor does it represent a research report. Securities may not be offered or sold in the United States absent registration with the US Securities and Exchange Commission or an exemption from registration under the US Securities Act of 1933, as amended. This document is only directed at professional investors who have experience of investing in emerging markets and the referenced investments are unlikely to be suitable for most private individuals. The referenced investments are speculative and include a high level of risk, and investors may not receive back the original amount of money that they invested. The value of investments can fall as well as rise, and you may get back less than what you originally invested. Where an investment is made in overseas currencies, changes in currency exchange rates may affect the value of your investment. Investments in emerging markets can be more volatile than in other more developed markets. Past performance is no guarantee of future performance, and the value of investments can go down as well as up. Please consult your financial and tax advisers if you are considering investing in any of the referenced investments. This document may contain certain forward-looking statements with respect to MidLincoln Research's strategies or expectations. Forward-looking statements speak only as of the date they are made, and MidLincoln Research assumes no duty to, and does not undertake to, update such forward-looking statements. This document may not be reproduced, distributed, transmitted, displayed, published or broadcast by any recipient for any purpose without the prior consent of MidLincoln Research. This document has been issued by MidLincoln Research.