

# The industry beneath the index.

Every sector ETF tells a flat story. XLE is energy. XLF is financials. XLK is technology. One ticker, one number, one beta to SPY.

But the *industry beneath the ETF* — the 4-digit slice the cap-weighted aggregate flattens — is rarely flat. Inside XLK there are semiconductors and software and IT services, each with its own typical  $\beta$  to the market and its own internal dispersion. Some periods, every industry inside XLK moves with the same  $\beta$ . Other periods, semiconductors detach. Software shrugs. The aggregate looks calm; the slices have rotated underneath.

**That rotation is the signal sector ETFs paper over.**

---

## The data product

For each EODHD 4-digit industry, on each trading day, at each cascade level (market / sector / subsector), RiskModels publishes four numbers:

Field	What it tells you
<code>beta_mean</code>	Log-mcap-weighted typical $\beta$ of the industry's members at that level.
<code>beta_variance</code>	Cross-sectional spread of $\beta$ <i>inside</i> the industry — high values mean the label is hiding real structure.
<code>n_companies</code>	Sample size at that level (and the universe gate it cleared).
<code>total_log_mcap_weight</code>	Cap-weighting denominator — the Vasicek peer mean's anchor.

The mean is what most allocators ask for: “what’s the typical chip-stock  $\beta$ ?” The **variance** is what most allocators *should* ask for: “how unified is the chip cohort right now?”

---

## Three things you couldn't see before

### 1. QUIET SECTOR ROTATION

A sector ETF can be flat while its industries rotate underneath. Track month-over-month change in `beta_mean` at `level='subsector'` across the 72 EODHD industries. When semiconductors' typical  $\beta$  jumps while the rest of tech holds steady, the *aggregate* XLK barely moves but the regime has rotated.

### 2. DISPERSION AS ALPHA SIGNAL

High `beta_variance` inside an industry flags structural variation hiding behind a single label. Some chip companies look like memory; some look like analog; some look like services. When the within-industry variance widens, single-name stock-picking is more rewarded than industry-level allocation. When it narrows, the cohort is moving together — index instruments do more of the work.

### 3. HEDGE-SIZING SANITY CHECKS

For an allocator hedging a sector exposure with the sector ETF, the cap-weighted industry `beta_mean` is the right notional to assume per dollar of sector exposure. Off the cuff most desks assume  $\beta=1.0$  across the sector. The actual  $\beta$  by industry in XLK varies. The industry panel lets you size the macro overlay correctly.

---

## How it works

The peer- $\beta$  statistics come from the **Vasicek shrinkage prior** the ERM3 engine uses internally to stabilize per-stock  $\beta$  estimates. The model already has to compute industry peer means and variances to shrink each stock's raw  $\beta$ ; the **industry-panel endpoint exposes those intermediate aggregates as a data product** instead of consuming them internally.

What used to be “what's the typical  $\beta$  for this industry?” — a question that required holding the whole universe in pandas — is now a single HTTP call returning  $\sim 72$  industries  $\times$  3 levels  $\times$  4 metrics. **Under 2 KB on the wire.**

```
from riskmodels import RiskModelsClient
client = RiskModelsClient.from_env()

# Industries by subsector dispersion, today
df = client.get_industry_panel(level="subsector", min_peers=20,
as_dataframe=True)
df.nlargest(10, "beta_variance")[["fs_industry_code", "beta_mean",
"beta_variance", "n_companies"]]
```

---

## Where it matters

- **Macro PMs** sizing sector overlays with the right per-dollar  $\beta$  per industry.
- **Sector rotators** tracking month-over-month change in `beta_mean` to surface quiet rotations the aggregate ETF flattens.
- **Single-name long/short books** using `beta_variance` as a dispersion-regime indicator — high within-industry variance is when stock-picking gets paid.
- **AI agents** answering “what’s the typical  $\beta$  for this slice?” without looping `get_metrics` over 30 names and reconstructing the aggregate by hand.

---

## The principle

Sector ETFs are summaries. Industries beneath them are structure. RiskModels publishes the structure as a first-class data product, not a derivation.

[Try the industry panel on the API →](#) · [Full endpoint docs →](#)