AlphaAgent

HOUSE RESEARCH DOCTRINE · GOVERNMENT OF CANADA CURVE

Government of Canada Curve — Relative Value

House doctrine for a specialist agent working the Canada curve.

How the house represents the curve, defines rich and cheap, expresses and sizes a relative-value view, assumes dealing cost, and audits its own work before any number is trusted.

PART ONE

Curve construction and rich/cheap

Principle

A yield curve is not a set of independent prices; it is one object, and the market prices time along it with considerable consistency. A single point can nonetheless drift away from where the rest of the curve says it should sit, usually for technical rather than fundamental reasons — an auction, an index extension, a hedging flow, a temporary shortage of a particular issue. **The house treats those deviations, not the level of yields, as the opportunity.** Everything therefore depends on having a defensible view of where each point *should* sit, and that view is a fitted curve. A residual is only as trustworthy as the fit that produced it.

The representation

The curve is represented by a parsimonious function of maturity, fitted to the observable benchmark tenors each day. Two families are acceptable: a parametric form whose parameters correspond directly to level, slope and curvature, and a smoothing spline where flexibility matters more than interpretation. A complementary, model-light view decomposes daily yield *changes* into their principal components, which for a government curve reliably recover the same three shapes and typically explain the overwhelming majority of daily variation. The house prefers the most parsimonious representation that fits: with only a handful of observable tenors, a flexible model will fit noise and manufacture residuals that are artefacts of the fit rather than facts about the market. **Whichever representation is chosen, the choice is written down and justified before any residual is computed.**

Stability is a precondition

A fitted curve that moves materially from one day to the next produces residuals that reflect the model's instability rather than the market's mispricing. Day-over-day parameter

stability and the fit error itself are therefore monitored as first-class diagnostics, not afterthoughts. **If the fit is not stable, the rich/cheap reading is not published.**

The residual and its sign

Richness and cheapness are defined as the deviation of an observed yield from the fitted curve at the same maturity. The sign convention is fixed and never inverted in reporting: a **positive residual means the observed yield sits above fair value, and the point is cheap**; a negative residual means the yield sits below fair value, and the point is rich. Because a cheap point is one whose price is low, buying cheapness means receiving a higher yield than the curve says that maturity deserves.

Conviction

A residual expressed in basis points is not yet a signal, because a two-basis-point deviation means something very different at a quiet point of the curve than at a volatile one. Conviction is expressed by standardising the residual against its own recent history — how unusual is today's reading relative to the way this point normally behaves. The lookback used to form that history is a stated parameter, the statistics are computed strictly from observations preceding the one being judged, and the threshold at which a reading becomes actionable is set in advance rather than chosen once the answer is visible.

Data integrity before modelling

Two properties of the published benchmark series are handled explicitly before anything is fitted. First, each series tracks whichever issue is currently the benchmark at that tenor, so when the market adopts a new benchmark the series changes underlying bond and a small discontinuity appears that is not a market move; the house either builds from day-to-day changes, which confines the contamination to a single observation, or adjusts explicitly at known switch dates, and documents which. Second, series drawn from different national calendars do not share holidays, so an explicit alignment rule is adopted and applied consistently. Where an independent official series covers the same quantity, the two are compared over their overlap and any disagreement is resolved before modelling begins, rather than discovered afterwards.

PART TWO

Relative value and backtest construction

Principle

Owning a cheap point outright is mostly a bet on the direction of yields, and direction will overwhelm the small relative mispricing the analysis identified. **The view is therefore always expressed market-neutral:** the cheap point is held against offsetting positions in its neighbours, so that broad moves cancel and only the targeted relationship drives the outcome. A relative-value position that makes money simply because yields fell has not been constructed properly.

The structure

The canonical expression of a view on a single point is the butterfly: the middle maturity, the belly, traded against the two surrounding maturities, the wings. The tracked quantity is the fly spread — twice the belly yield less the two wings — which measures how far the belly sits from the straight line between its neighbours, and is therefore a direct reading of curvature at that part of the curve. A high fly spread means the belly is cheap to its wings; a low one means it is rich. Owning the belly profits as the fly spread falls back toward its average. The primary structure is stated in advance, and a second structure at a different part of the curve is carried as a check that a result is a property of the curve rather than of one tenor.

Weighting the legs

How much of each wing is traded against the belly determines what the position is actually exposed to. Matching face amounts leaves the position exposed to the direction of yields, because a long-maturity dollar is far more rate-sensitive than a short-maturity one. **The house sizes by dollar sensitivity rather than face value**, so the wings together carry the same sensitivity as the belly and a parallel move in all yields nets to approximately nothing. Where the view is specifically about curvature, weights may instead be chosen so the position is neutral to both the level and the slope factors, leaving exposure to curvature alone; that is cleaner, has more moving parts, and the choice is stated either way.

The cost assumption

Public sources carry no quoted bid/ask and no depth, so dealing cost cannot be measured and must be assumed. This is treated as a first-order modelling decision rather than a footnote. A butterfly trades three legs and a mean-reversion rule re-trades often, so cost compounds directly against the edge. The house assumes a conservative fixed cost per leg expressed in yield terms, charges it on **every leg of every entry, exit and rebalance**, and reports the headline result beside versions computed at twice and three times that assumption. **A result that does not survive a doubling of assumed cost is reported as a failure, not as a result with a caveat.**

Simulation discipline

A historical simulation is evidence only to the extent that it could not have known what it was about to find. Signals may use only information available on the day they are formed; positions taken on a signal execute at the following day's prices; and any rolling statistic excludes the observation it is being used to judge. Any parameter that is fitted is estimated on past data only and applied forward to data the estimation never saw, rolling through the sample. Nothing is reported as a single tuned configuration: results are shown across a grid of reasonable lookbacks, thresholds and cost assumptions, because a genuine effect survives its neighbourhood and an artefact does not. Turnover and average holding period are reported alongside return, because a result that depends on trading constantly is a different proposition from one that does not.

PART THREE

Stress, validation and adversarial audit

Principle

An average describes ordinary conditions; risk is decided in the conditions that are not ordinary. A relative-value position that behaves well on a typical day can behave very badly in the few weeks that determine a year. **The house therefore requires the stress profile to be computed rather than asserted**, and computed from the same data and the same code as everything else, so a stress number carries the same auditability as a headline number.

Regime conditioning

Curve behaviour differs profoundly with what the central bank is doing, so no result is reported as a single full-sample figure. Each observation is labelled by the phase of the policy cycle — tightening, easing, or on hold — derived from the published policy rate, and performance is reported by phase. Moves are additionally described in the standard vocabulary, combining the direction of yields with the change in shape, so the conditions under which the strategy earns or loses are named rather than averaged away. A strategy that earns only while the central bank is on hold is a materially weaker proposition than one that survives a tightening cycle, and the doctrine requires that distinction to be visible.

Historical replay

The position is run through the episodes that actually tested this market, day by day, reporting the path rather than only the endpoint: the collapse in growth expectations and flight to safety of the global financial crisis; the weeks in March 2020 when even government bond markets malfunctioned and liquidity in anything other than the most current issue evaporated; and the fastest tightening cycle in decades, which drove a deeply inverted curve and presented mean-reversion logic with a persistently directional market. **The question asked of each episode is not only what was lost, but whether the signal remained coherent** while it was being lost.

Hypothetical shocks

Alongside the replays, a defined set of instantaneous curve moves is applied to current positioning and the immediate result reported: a parallel shift in both directions, a flattening driven by the short end, a steepening driven by falling short yields, and a sharp long-end rally of the kind seen on panic days. The parallel shift doubles as a construction check — a position sized neutral to the direction of yields should be close to unaffected by it, and confirming that is part of validating the hedge rather than a formality.

The checks

A result is reported only when every one of the following holds.

- **Look-ahead.** No signal uses information that did not exist when it was formed, no rolling statistic includes the observation it judges, and no position is filled at the price

that generated it.

- **Estimation discipline.** Every fitted parameter was estimated on data preceding the period it is applied to and rolled forward; nothing was selected after the outcome was visible.
- **Fit integrity.** The curve fit was stable across the period, and the residuals being traded are not artefacts of an unstable or over-flexible representation.
- **Cost realism.** A conservative cost was charged on every leg and every rebalance, and the conclusion was re-tested at twice and three times that assumption.
- **Regime robustness.** Performance was decomposed by policy phase, and any dependence on a single regime is stated rather than buried in a full-sample average.
- **Neighbourhood stability.** The result holds across a range of reasonable lookbacks and thresholds; a result that exists only at one parameter setting is treated as an artefact.
- **Data provenance.** Every figure resolves to a named source series and a dated retrieval, so the work can be re-run from scratch and reach the same numbers.

The verdict

Acceptance thresholds are stated before the analysis runs, not after: the return per unit of risk that would make the strategy worth carrying, the statistical bar that separates an edge from a lucky stretch, the requirement that performance be positive across more than one policy phase, and the limit on what may be lost in the worst replay or shock. For each threshold the agent judges pass or fail, states the specific number, and names the assumption most responsible where a threshold fails. **A clean, well-evidenced failure against a stated threshold is a complete and successful piece of research**, and is reported as such rather than reworked until it passes.