


From Tools to Teammates: Integrating AI Agents into Data & Analytics

How AI agents participate across the analytics lifecycle

Anusha Dwivedula,
Director of Product,
Analytics, Morningstar


Josh Charney, Associate
Director of Quantitative
Research, Analytics,
Morningstar.



We are the global leader for investor-first insights, tools, and investment strategies, trusted by market participants worldwide. Our system of capabilities, connected through data and insights, is driven by our mission to empower investor success.

We'll follow one analytic from idea to client — not to show five AI demos, but to look at what happens when **agents participate in the work.**

Explore → Understand → Build → Operate → Deliver



Foundation: Data · Context · Tools · Governance · Evals

Client questions become new research

Explore: designing the investigation

The hard part isn't producing the answer. **It's designing the investigation.**



Before

- Searches documents and prior research
- Reviews existing methodologies
- Hunts for candidate datasets
- Forms hypotheses
- Sets the plan



With agents

- **The human** frames the problem
- **The agent** gathers context and surfaces methodological options, assumptions and open questions
- The plan is settled together
- The researcher owns whether it is sound and defensible

Understand: fitness for purpose

Finding data is a search problem. **Deciding whether it's the right data is an analytical problem.**



The data platform

- Governed datasets and definitions
- Semantic layer and lineage
- Entitlement-based access



The agent

- Profiles coverage and missingness
- Flags outliers and distribution shifts
- Tests for hidden bias



The human

- Approves proxies and exclusions
- Adapts methodology to data reality
- Signs off for production

From methodology to production code

The leap isn't that AI can write Python. It's giving AI enough **structure and context to turn analytical intent into reliable implementation.**

Methodology requirement	Implementation	Test
Minimum 36 months of history	Eligibility function	Boundary tests at 35, 36 and 37 months
Cap extreme observations	Outlier-processing module	Extreme-value test
Use category-relative comparison	Peer-group calculation	Category membership test

What monitoring tells you today

ALERT 04:12 UTC

Calculation variance exceeded threshold.

That is everything traditional monitoring tells you.

Not why. Not who is affected. Not what to do next.

From alert to evidence-backed diagnosis

The agent's response

The change affects **2.8% of portfolios**, concentrated in three fixed-income categories. Most likely cause: a classification update in today's holdings load. The calculation service is healthy. Reprocessing with the prior classification restores expected results.

Human approval required before rollback.

Triage time drop to hours instead of days

Bounded autonomy



The first win isn't autonomous remediation. **It's collapsing the time from detection to diagnosis.**

From field names to client intent

“Which funds in this portfolio have the highest downside risk, and what's driving it?”



Before

- Know which product holds the analytic
- Know which field to use
- Know how to interpret it
- Know what limitations apply



Now

- Interprets client intent
- Selects approved analytics and retrieves entitled data
- Explains the result and cites methodology
- Surfaces limitations

Grounded in governed data, approved analytics, published methodologies and client entitlements — **it does not invent a new analytic.**

Data Product -> API -> Agentic Capability

APIs distribute data to applications.

Agents distribute analytical capability into decisions.

The screenshot displays the Morningstar Data+Analytics web interface. At the top, the navigation bar includes the Morningstar logo, 'Data+Analytics', and a 'Morningstar AI Solutions' dropdown menu. A 'Talk to us' button is located in the top right corner. Below the navigation bar, the breadcrumb trail reads 'Home / Products / Morningstar MCP Server'.


The main content area features a section titled 'Connect trusted data to your AI workflows with the Morningstar MCP Server'. This section includes a paragraph: 'Ground applications like Claude, ChatGPT, and Copilot in trusted Morningstar data and research.' and another paragraph: 'From investment selection to portfolio reporting, the Morningstar MCP Server gives you access to governed, AI-ready data and analytics—right where you need it.' Below this text are two buttons: 'Request a free trial' and 'Talk to sales'.

To the right of the text, there is a vertical list of tools with 'Result' indicators: 'Morningstar Analyst Research Tool', 'Morningstar Portfolio Analysis Tool', and 'Morningstar Fund Holding Tool'. Further right, a panel displays logos for AI agents: ChatGPT, Claude, Microsoft 365 Copilot, and Perplexity.

The bottom half of the interface shows a 'Portfolio X-Ray' dashboard. It includes a 'New Chat' button, a 'Projects' list, and a 'Recents' list. The main dashboard area displays 'Asset Allocation - Net Exposure VS S&P 500' with a bar chart and four data points: Non-US equity (31.0%), Bonds (86.5%), US Equity (0.8%), and Cash (-6.7%). A legend at the top of the chart area lists: VWIGX 35%, PIMIX 25%, CRDIX 20%, MQY 10%, and Cash 10%.

Teammates have roles. Agents should too.

Not one super-agent — specialized roles with explicit hand-offs and human decision points.



medalist-rating-analyzer

Search → fetch and normalize across three MCP tools → format the report. The user selects the matching fund before anything is fetched.

medalist-rating-methodology

Formulas, scales and framework behind the rating. Explicitness gate: answer only what the source states. Fund-specific lookups hand off to the analyzer.

USE CASE

What is the Medalist Rating?

175K · 360K+


traditional funds/ETFs · all vehicles globally
rated as of Mar 31, 2026

3,200+

fund strategies
qualitatively rated by analysts

130

global researchers
manager research team



HOW IT'S ASSIGNED

Analyst

Qualitative pillar ratings


Algorithm

Quantitative rules-based model

The New Simplified Medalist Rating Input Weights

	Active Vehicles	Passive Vehicles
Pillar Ratings Morningstar reviews three pillars to reinforce a consistent, long-term investment process.	70% 45% People Rating 45% Process Rating 10% Parent Rating	60% 10% People Rating 80% Process Rating 10% Parent Rating
	+	+
Price Score The Medalist Rating Price Score reflects an investment vehicle's fee in relation to its Morningstar Category peers.	30%	40%
	=	=
Medalist Rating	★ Gold ★ Silver ★ Bronze Neutral Negative	★ Gold ★ Silver ★ Bronze Neutral Negative

How a question becomes an answer



Specialized skills, explicit handoffs

Not one super-agent — two governed skills with clear roles

MORNINGSTAR MCP SERVER

medalist-rating-analyzer

Search → fetch and normalize across three MCP tools → format the report. The user selects the matching fund before anything is fetched.

1 RESOLVE

Fund name / ticker / ID → correct investment



HUMAN CHECKPOINT — identifier ambiguous?

Yes → ask the user to clarify. No → continue.



2 FETCH

Identity · research · data + history



3 PRESENT

Rating · pillar scores · history

medalist-rating-methodology

Formulas, scales and framework behind the rating. Explicitness gate: answer only what the source states. Fund-specific lookups hand off to the analyzer.

Formulas

Rating scales

Caps and buffers

Eligibility rules

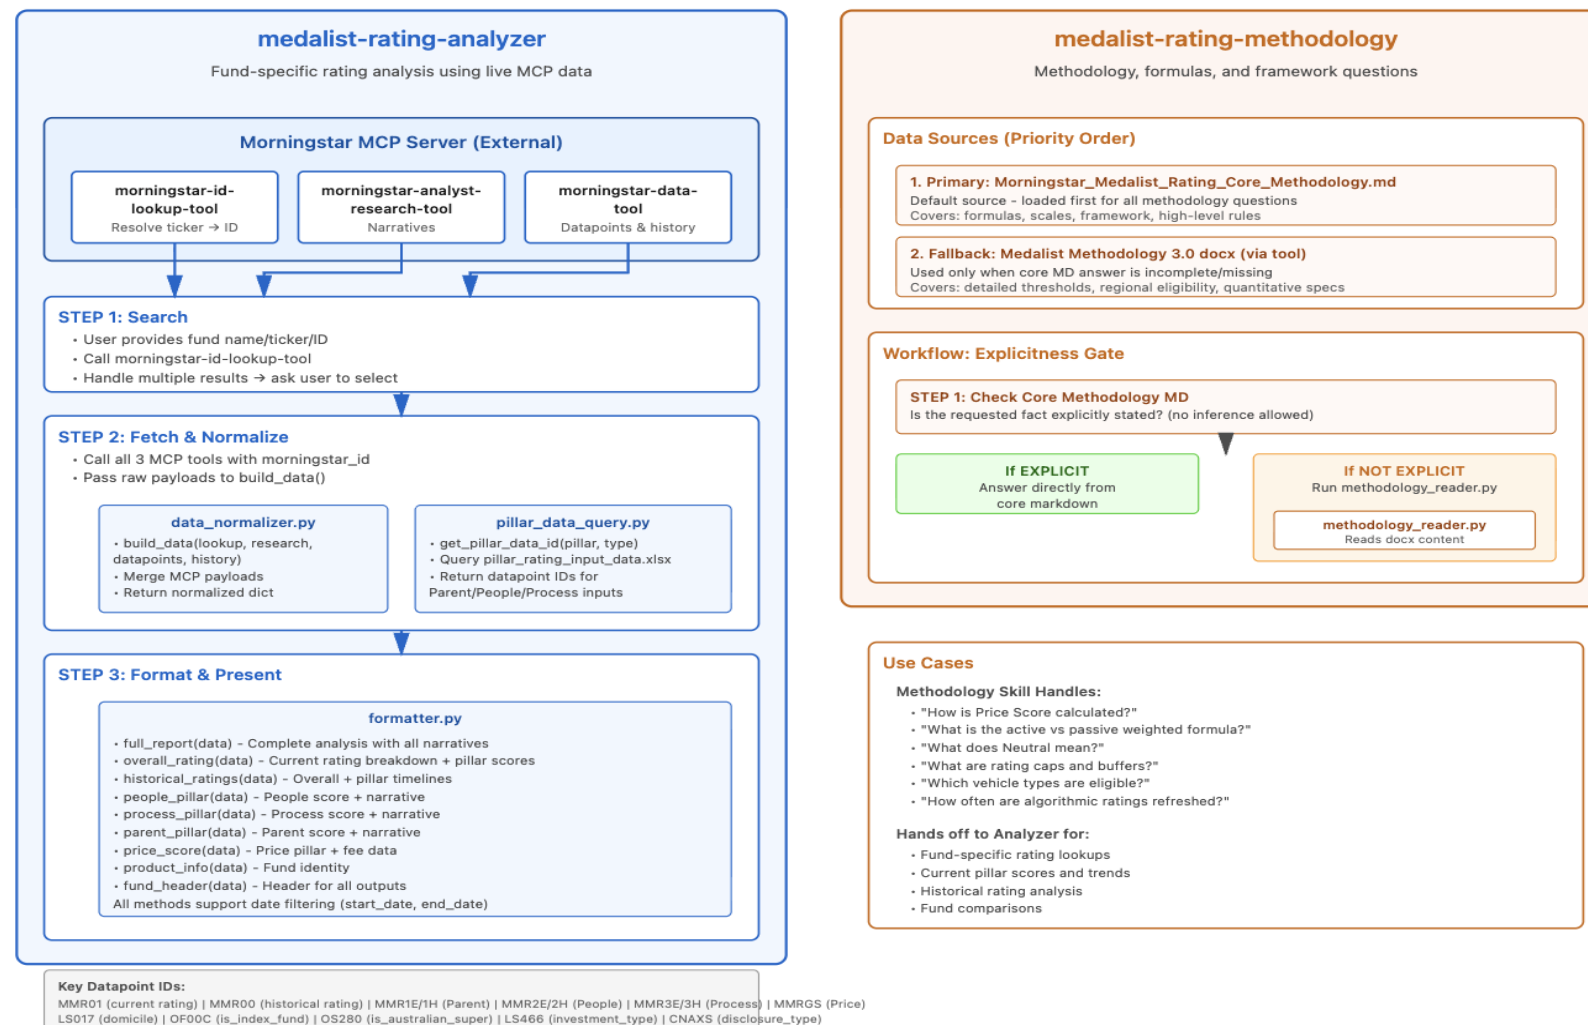
Refresh cadence

← FUND-SPECIFIC QUESTION → HAND OFF TO THE ANALYZER

Specialized roles. Explicit contract. Human checkpoint. Clear ownership.

Morningstar Medalist Rating Skills Architecture

Component structure and data flow for medalist-rating-analyzer and medalist-rating-methodology



What makes this trustworthy



Governed data

Trusted,
lineage-traced,
entitlement-
aware



Semantic context

Definitions,
methodologies
and meaning a
machine can
read



Tools

The
Morningstar
MCP Server:
real systems
agents can act



Governance

The explicitness
gate: answer
only what the
source states



Evaluation

Benchmarks,
failure analysis,
domain experts
in the loop

The organizations best positioned for this aren't the ones with the best model. **They're the ones whose data, definitions, methodologies and lineage are accessible to machines.**

At what point did these stop looking like tools?

Not when they became conversational. Not when they could generate code.

They started looking like teammates when they could take a goal, use our data and tools, do a bounded piece of work, evaluate the result — and **know when to bring a human back in**.

THE ASK

Make our data, definitions, methodologies and lineage machine-accessible. That is the prerequisite for everything shown here.

THANK YOU!

Get in touch



Anusha Dwivedula
Director of Product,
Analytics, Morningstar



Josh Charney
Associate Director of Quantitative
Research,
Analytics, Morningstar.