MCP SERVER | NO CODE | CLOUD HOSTED

Bayesian A/B Testing Calculator MCP for AI Agents

Calculating True Conversion Probability and Risk in CRO

Bayesian A/B Testing Calculator uses advanced statistical methods to evaluate website variant performance. Stop relying on simple p-values; this MCP quantifies conversion probability and expected loss with Bayesian inference, telling you exactly how confident you should be in a winner.

A+ Quality Score 100/100

bayesian | ab-testing | conversion-rate | statistical-inference | data-analysis

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/bayesian-ab-testing-calculator

SCAN TO VIEW

The connectivity layer between AI and the world's software.

Vinkius sits between AI and every application. All communication passes through Vinkius Cloud via the Model Context Protocol (MCP) — with governance, observability, and security at every layer.

Your AI Connections Run Through Vinkius Cloud

The world's largest managed MCP catalog

Vinkius is the connectivity layer where AI connects to the software your business already runs. We handle the hosting, the security, the credentials, the uptime — you get agents that actually do things.

We operate the world's largest managed MCP catalog. Major SaaS platforms, CRMs, databases, and cloud providers — running, monitored, production-ready. This MCP server is hosted and maintained by the Vinkius Cloud for AI Agents.

> *The agent doesn't manage credentials, doesn't manage uptime, doesn't manage security. Vinkius does.*
>
> — Architecture principle

Four Pillars of the Vinkius Runtime

01 — Security by design

Credentials stay encrypted at rest via AES-256. The AI agent never touches raw keys — they're injected into a sandboxed V8 isolate at runtime. Actions are logged, and connections have an emergency kill switch.

02 — Built on MCP Fusion

This MCP server was built with **MCP Fusion**, the open-source framework (Apache 2.0) that powers the entire Vinkius catalog. Schema-as-firewall strips undeclared fields, compiled PII redaction runs at zero overhead, and cryptographic lockfiles produce git-diffable audit trails.

03 — Deterministic observability

Eight immutable metrics per endpoint: request volume, p95 latency, error rate, active connections, cost attribution. A live payload feed logs every tool call with mutation detection.

04 — Autonomous operations

Servers are deployed, monitored, and patched autonomously. New capabilities and security patches ship weekly. Zero-downtime deployments ensure continuous availability across all managed MCP servers.

AES-256	Ed25519	24h TTL	V8 Isolate
Encryption at rest	PKI vault signatures	Ephemeral session keys	Sandboxed execution

One Token. Instant Access.

Every MCP server on Vinkius is accessed through a **Connection Token**. Tokens are generated in the cloud dashboard and produce a unique MCP endpoint URL. Paste this URL into any MCP-compatible client — no SDK required.

A single token can serve **multiple AI clients simultaneously**, or you can issue separate tokens per client for granular access control. Each token tracks its own request count, last activity timestamp, and can be individually enabled or revoked.

MCP ENDPOINT

```
https://edge.vinkius.com/{token}/mcp
```

Claude · Cursor · VS Code · Windsurf · Grok · Gemini

Security Is the Architecture

Security in Vinkius is not a feature — it's the foundation of the runtime. The gateway enforces multiple independent protection layers between AI agents and third-party APIs.

01 — Ed25519 PKI Vault

Every workspace has an Ed25519 Master Key. Session keys are generated ephemerally (24h TTL) and signed by the Master Key. Credentials never leave the vault boundary.

02 — V8 Isolate Sandboxing

Tool code runs inside isolated-vm V8 isolates with 64 MB memory caps and per-request timeouts. No filesystem access, no network access except through the SSRF-guarded fetch bridge.

03 — SSRF Guard

All outbound HTTP requests are DNS-resolved and validated before execution. Private IP ranges (10.x, 172.16-31.x, 192.168.x, AWS metadata 169.254.x) are blocked at the network layer.

04 — DLP & PII Redaction

A ResponseGuard pipeline intercepts every tool response. Configurable redaction patterns strip sensitive fields (emails, SSNs, card numbers) before data reaches the AI agent.

05 — Cryptographic Audit Trail

Every request is signed into a SHA-256 hash chain with Ed25519 signatures. Events form a tamper-proof, SIEM-exportable forensic record.

06 — Honeytoken Trap System

Phantom credentials are injected into isolated environments. If a honeytoken is used outside Vinkius infrastructure, the server is quarantined instantly.

Emergency Kill Switch

EU AI Act Art. 14(1) Compliant

The kill switch is an **emergency halt** mechanism — not a simple toggle. When triggered, it executes three actions atomically:

01 — Server deactivated

The MCP server is immediately taken offline across the entire cluster.

02 — All tokens revoked

Every connection token is invalidated. Total lockout — reconnection blocked until new tokens are issued.

03 — WebSocket connections killed

Active connections terminated via Redis pubsub broadcast. Propagates to every runtime node in the cluster.

Full Visibility. Zero Guesswork.

The Vinkius cloud dashboard includes a full MCP Governance suite — real-time analytics and security controls for production AI operations.

Control Plane

KPI dashboard with request volume, latency, success rate, token consumption, and AI-generated operational briefings.

Agent Activity

Which AI clients are connecting, how often, and what they're doing — real-time session tracking.

FinOps

Cost tracking per tool, payload compression savings, budget optimization signals, and consumption trends.

Tool Health

Slowest and most error-prone tools, with actionable root-cause insights and performance baselines.

Firewall & DLP

PII redaction activity, sensitive data protection counters, and security event timeline.

Incident Log

Error trends, failure rates, status-code breakdowns, and forensic audit trail access.

> Get started at **cloud.vinkius.com** — connect your AI agent in under 60 seconds.

Bayesian A/B Testing Calculator MCP

4 tools available
Cloud-hosted on Vinkius

When running A/B tests, simply checking the p-value doesn't give you the full picture. This MCP provides a powerful statistical engine that moves beyond basic significance testing. It uses the Beta-Bernoulli relationship to calculate the actual probability of one variant beating another. Instead of just flagging a difference, your agent tells you how much risk you take by making a wrong decision using tools like `calculate_expected_loss`. You can also determine exactly what uplift Variant B provides over Variant A with `calculate_expected_uplift`, giving you clear numbers for product prioritization and resource allocation. Once the math is done, the MCP helps guide your next steps through `evaluate_decision_recommendation` or confirm confidence levels using `calculate_superiority_probability`. By connecting this to Vinkius, you give your AI client instant access to sophisticated analytics that most internal tools just can't match.

Core Capabilities

01 — Determine Winner Confidence

Calculates the precise probability that one variant significantly outperforms another.

02 — Quantify Decision Risk

Measures the expected loss associated with choosing either variant before testing is complete.

03 — Estimate Performance Gains

Projects the anticipated uplift in conversion rate for one variant over a baseline.

04 — Generate Actionable Next Steps

Provides clear, data-driven recommendations based on your required confidence threshold.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/bayesian-ab-testing-calculator** — connect your AI agent in three steps.

01 Feed the MCP raw A/B test performance metrics—specifically conversion counts and total visitors for both variants.

02 The engine processes this data using Bayesian inference to generate probability distributions, moving past simple statistical significance.

03 Your AI agent synthesizes these results into clear numbers, telling you the likelihood of superiority or the expected financial loss if you choose incorrectly.

The bottom line is that it converts raw traffic logs and conversion counts directly into actionable business probabilities, eliminating guesswork from product decisions.

Built For

Product managers and data analysts who are sick of vague statistical reports and need hard numbers to justify expensive design changes. If your team spends hours arguing over whether a metric is 'significant enough,' this MCP is for you.

Conversion Rate Optimization (CRO) Manager

Uses the MCP to compare multiple landing page versions, calculating expected uplift to prove ROI before launch.

Product Data Analyst

Runs complex simulations on user flow data, using Bayesian analysis to determine which feature change minimizes risk and maximizes conversion probability.

Marketing Director

Needs quick, reliable answers about campaign performance, relying on the MCP to provide clear decision recommendations based on defined business thresholds.

What Changes When You Connect

01 Moves beyond unreliable p-values. Instead of just knowing if a difference exists, you know the *probability* of that difference being real.

02 Quantify risk immediately. Use `calculate_expected_loss` to see the potential cost of making a bad product decision before you launch anything.

03 Pinpoint concrete gains with `calculate_expected_uplift`. You get a single, clear number showing how much better one variant is predicted to be.

04 Get direct guidance. The MCP uses `evaluate_decision_recommendation` to tell you what to do next, based on the confidence level your business requires.

05 Saves days of manual analysis. Instead of running complex scripts in R or Python, your agent handles the entire statistical pipeline instantly.

Real-World Applications

Landing Page Optimization Failure

A CRO Manager runs two versions (A and B) of a sign-up page. They ask their agent to calculate the superiority probability, confirming that B has a 98% chance of being better than A, allowing them to greenlight the rollout immediately.

Feature Rollout Risk Assessment

A Product Analyst is debating between two pricing models. They use the expected loss calculator to determine that choosing the wrong model carries a $50k risk, forcing the team to prioritize testing high-impact changes first.

Identifying Minimum Viable Changes

A Marketing Director needs proof that a small copy change is worth the effort. They use the expected uplift calculation and see a projected 4% gain, proving the investment will pay off.

Patterns to Avoid

Relying solely on p-values

✗ AVOID

Assuming that since the p-value is <0.05, Variant B *must* be better and implementing it immediately without understanding the probability distribution.

✓ INSTEAD

Always calculate the superiority probability first using `calculate_superiority_probability`. This gives you a much more reliable chance of success than simple significance testing.

Ignoring potential losses

✗ AVOID

Launching a new feature because it looks 'better' in a dashboard, without accounting for the financial risk if that assumption is wrong.

✓ INSTEAD

Before committing resources, use `calculate_expected_loss`. This tool forces you to quantify exactly what you stand to lose by picking the suboptimal variant.

Confusing correlation with causation

✗ AVOID

Seeing that Variant B performed better last week and assuming it will always outperform A, without recalculating based on current data metrics.

✓ INSTEAD

Run a fresh analysis using all the tools. The `evaluate_decision_recommendation` tool forces you to check your findings against a specific business confidence threshold.

The Right Fit

Use this MCP if your product decisions hinge on statistical proof, not gut feeling. You need to know not just *if* there's a difference (p-value), but the true probability of superiority and the expected financial risk involved. Don't use it if you simply need to track metrics or segment users; that requires a dedicated analytics dashboard. If all you have is basic comparison data, run `calculate_superiority_probability` first. Only rely on an action when both the superior probability is high AND the calculated loss for making the wrong choice is manageable.

Bayesian A/B Testing Calculator: Solving Conversion Rate Uncertainty

Today, running a simple A/B test means exporting data into a spreadsheet and calculating p-values. This process is slow, tedious, and worse, it tells you half the story. You get a binary 'significant' or 'not significant,' but that doesn't help your product roadmap.

With this MCP, your agent takes raw conversion counts and visitors and instantly calculates true probability using Bayesian methods. You stop guessing about your conversions; you start making decisions based on quantified certainty.

Bayesian A/B Testing Calculator: Quantifying Product Risk

Manual testing often stops at identifying the winner, leaving the team blind to the cost of failure. You waste time debating which minor change is worth rolling out because you haven't quantified the actual risk.

The MCP calculates your expected loss and uplift in one go. It shifts the focus from 'which variant wins?' to 'what are we willing to risk, and how much will this win make us?'

Bayesian A/B Testing Calculator: 4 Tools for Conversion Rate Analysis

Use these four tools to calculate superiority probability, measure expected loss, project uplift, and get final recommendations from your A/B test data.

#	TOOL	DESCRIPTION
01	`evaluate_decision_recommendation`	Gives an actionable recommendation for a winner or loser based on your specific confidence threshold percentage.
02	`calculate_superiority_probability`	Calculates the precise probability that one variant is better than another, moving beyond traditional statistical tests.
03	`calculate_expected_loss`	Determines the expected loss in revenue or conversions if you choose the wrong variant for your product.
04	`calculate_expected_uplift`	Calculates how much better one variant is predicted to perform compared to a baseline version.

See It in Action

Real prompts you can use once this MCP is connected to your AI agent through Vinkius Cloud.

U We tested a new checkout flow (B) vs old one (A). A had 80 conversions / 2000 visitors; B had 110 / 2000. What's the chance B is actually better?

Superiority Probability Calculation

The probability that Variant B outperforms Variant A is approximately **96.3%**.

This means we can be very confident in this result.

- **Next Step:** Calculate expected uplift to confirm value.

U If I choose the current flow (A) over the new one (B), what's the estimated loss? A: 50/1000, B: 60/1000.

Expected Loss Analysis

Choosing Variant A results in an expected loss of **0.4%** compared to Variant B.

This quantifies the risk for leadership: a decision against the better variant costs us measurable conversions.

U We need a final recommendation for A (20/200) and B (35/200). Our threshold is 98% confidence.

Decision Recommendation

Based on the 98% confidence threshold, the decision is **RECOMMENDED: Variant B**.

The probability of superiority for B meets your minimum required certainty level. Implement this change immediately.

Frequently Asked Questions

01 **How does the Bayesian A/B Testing Calculator improve on standard statistical tests?**

It moves beyond simple p-values by calculating a true probability distribution. Instead of just saying 'there is a difference,' it tells you how confident your team should be in that finding, providing much clearer direction.

02 **Can I use this MCP to calculate the potential financial risk of poor A/B test results?**

Yes. You can run the expected loss calculation to quantify exactly what you stand to lose if your team makes a decision based on incomplete or misleading data.

03 **Do I need coding knowledge to use the Bayesian A/B Testing Calculator MCP?**

No. Your AI agent handles all the complex statistical math behind the scenes. You just provide the raw conversion counts and visitor numbers, and it returns clear percentages.

04 **What if I have more than two variants to test? Does the calculator handle that?**

The MCP is designed for comparing pairs of variants (A vs. B). You can run multiple comparisons sequentially to build a comprehensive picture of performance across all versions.

05 **How do I know if my results are 'good enough' to launch with the Bayesian A/B Testing Calculator?**

The tool provides an explicit decision recommendation. You set your confidence threshold, and it tells you precisely whether the data meets that business standard for a go-ahead.

Go Live in 60 Seconds

Get your connection token from **cloud.vinkius.com**, then paste the endpoint URL into any MCP-compatible client.

YOUR MCP ENDPOINT

```
https://edge.vinkius.com/[TOKEN]/mcp
```

CLIENT	WHERE TO CONFIGURE
Claude AI	Profile → Customize → Connectors → "+" → Add custom connector → Paste endpoint
Cursor	Settings → Features → MCP Servers → "+ Add New MCP Server" → Type: SSE → Paste endpoint
VS Code	Ctrl/Cmd+Shift+P → "MCP: Add Server" → add `"bayesian-ab-testing-calculator": { "url": "…" }`
Windsurf	MCP Settings → `mcp_settings.json` → Add endpoint URL
ChatGPT	Settings → Tools & plugins → Add MCP server → Paste endpoint
Gemini	Extensions → Add MCP Server → Paste endpoint URL

ASK AN AI ABOUT THIS

Let your preferred AI explain this MCP server

- Ask ChatGPT
- Ask Claude
- Ask Perplexity
- Ask Gemini
- Ask Grok

READY TO CONNECT

Bayesian A/B Testing Calculator is live on Vinkius Cloud.

Get your connection token, paste it into your AI agent, and start building. No SDK. No deployment. Just results.

Start at cloud.vinkius.com →

vinkius.com · support@vinkius.com

INDEPENDENT PLATFORM DISCLAIMER

Vinkius is an independent platform and is not affiliated with, endorsed by, sponsored by, verified by, or otherwise authorized by Bayesian A/B Testing Calculator. All third-party trademarks, logos, and brand names are the property of their respective owners. Their use in this document is strictly for informational purposes to identify service compatibility and interoperability.

DOCUMENT INFORMATION

Generated	June 2026
MCP Server	Bayesian A/B Testing Calculator MCP
Server ID	`019f11d5-9c80-7018-9ffa-97cd91595552`
Platform	Vinkius Cloud for AI Agents
Endpoint	`https://edge.vinkius.com/{token}/mcp`

LICENSE & USAGE

This document is generated automatically by the Vinkius PDF Engine. Content reflects the MCP server configuration at the time of generation and may change as updates are deployed. For the most current information, visit vinkius.com/mcp/bayesian-ab-testing-calculator.