

Aleksei Beliankin-Bauer

Researcher · Writer · Builder of autonomous research systems

abeliankin@legionfarm.com

ab7.ai · linkedin.com/in/aleksei-beliankin-bauer · x.com/ABeliankin · medium.com/@abeliankin

Builds and operates autonomous AI research systems in production, and applies the same evidence discipline to primary-source investigation — separating what is established from what is inferred from what someone wants to be true.

RESEARCH AND WRITING

Primary-source investigation — Dr. John E. Mack

August 2026

Original archival and records research into the Harvard psychiatrist's 1994–95 investigation and its funding. Established from IRS Business Master File records that the Laurance S. Rockefeller Fund now filing with the IRS received its exemption ruling in **June 2006** — two years after Rockefeller's death and thirteen years after the grant it is universally assumed to have made. The foundation every researcher would naturally check therefore *cannot* be the source, and a diligent journalist checking it draws a false negative from authentic documents.

Traced the most-repeated figure in thirty years of writing on the subject — "\$1M over four years" — to a source that says two.

Evidence-graded research corpus

August 2026

Multi-source dossiers in which every claim carries an explicit class — fact, reported, inference, speculation, or *missing-evidence* — with source, date, quotation, and the specific observation that would disprove it. The standing rule, quoted from the working documents: every number carries a source; anything unsourced is written as missing-evidence.

Deep-research fleet

August 2026

Parallel multi-agent research runs against primary sources, producing sourced reports in roughly twenty-five minutes of wall-clock time.

Profile-enrichment research pipeline — AGI House

June – November 2025

Designed and ran the applicant research and scoring pipeline behind an attendee-selection console — multi-role score vectors, a dated event briefing, longitudinal founder-impact tracking. **13,619** of **13,733** records carry the deployed scorer's completion flag, against **\$1,861.47** of provider spend summed from the transaction log across nine services.

An earlier draft said "13,733 profiles at forty-plus tool calls each." Both halves were wrong — 13,733 is the record count, not the scored count, and the call quota was a design target whose own audit field appears nowhere in the 450 MB of output. [Retraction, and how it was caught, on the verification record.](#)

OPERATING RECORD

Founder and CEO — LegionFarm / LFCarry

2016 – present

Founded the gaming-services marketplace and still run it; **Y Combinator**, **Winter 2020**. Company and YC profiles report cumulative lifetime scale of over **\$45M** in sales and **250,000+** customers served. Raised a ~\$1.7M seed in 2020 and a **\$6M** Series A in 2021, both covered by TechCrunch. Still CEO; the company now runs on three people and the agent fleet below.

Autonomous agent systems in production

2026 – present

Two production services running against a live marketplace — a buyer-facing conversational sales agent and a Slack-triggered content-engineering agent with deterministic safety guards — plus a governed multi-agent operating framework.

RECOGNITION

Forbes Russia, "**30 Under 30**," **2021** — Entrepreneurs.

TECHNICAL

Multi-agent orchestration (Cursor SDK, Codex) · Firecrawl and authenticated-browser research automation · Python · Postgres, MySQL · production deployment and observability.

EDUCATION

Novosibirsk State University — geology and geophysics. Left to build.