

PART 2 — PROOFS

Each Use Case, Run Against Live Data

— . —

PERSONA A — DESIGNERS

Three mini-reports. Each one runs the use case against the live BigQuery gold-tier views and reports the structured output. Where data is sparse or the anchoring event sits outside our collection window, that limit is named in the open. Persona B and C tranches follow.

The Unearthed Arcana Sentiment Auditor

The question. When WoTC drops a UA playtest packet, what does the community actually think — broken down by structured narrative type, not by raw sentiment polarity?

What we ran. WoTC has not shipped a fresh UA packet inside our v5 forum-body collection window (Apr 2026). So we ran the same classifier against the IP-reception corpus — every confirmed forum thread classified by the v5 narrative-type tagger across EN World, RPG.net, and Giant in the Playground. The classifier is the product. The input is interchangeable: UA packet today, IP crossover yesterday, rules-revision tomorrow.

```
SELECT ip_name, confirmed_about_ip_count, attitude_avg, cash_grab_count,
tone_mismatch_count, not_dnd_count, pandering_count, system_design_critique_count,
worldbuilding_endorsement_count, sample_backlash_evidence FROM
gold_data.forum_presence_proxy WHERE confirmed_about_ip_count >= 5 ORDER BY
(cash_grab_count + tone_mismatch_count + not_dnd_count + pandering_count) DESC LIMIT 5;
```

What came back. Top-5 backlash-tagged IPs across the three forums. Structured columns are confirmed-thread counts — every cell is a real forum post the AI Bouncer matched into that narrative class:

IP	Threads	Tone-mismatch	Cash-grab	Pandering	Design-critique	Endorsement	Attitude avg
Goblin Slayer	10	4	1	1	6	2	0.04
Stranger Things	10	1	2	2	1	1	0.36
Baldur's Gate 3	10	0	2	1	3	1	0.32
Welcome to Night Vale	7	2	0	0	2	0	0.18
The Boys	10	2	0	0	2	0	0.20

Source: `gold_data.forum_presence_proxy`, snapshot 2026-04-30, 142 IPs scanned, 24 above the 5-thread minimum.

What it tells us. The classifier produces six narrative tags per thread, not a single positive/negative score. That structure is what makes it useful for a Day-1 UA-packet read: "4 tone_mismatch flags out of 10 threads" tells the design team something specific about WHY a packet is divisive (the flavor reads wrong), separable from "6 system_design_critique flags" (the math is being argued). For Goblin Slayer, the triangulation shows up cleanly — the same IP is flagged tone_mismatch on three forums independently, which is the structural signature an OGL-tier event would produce on Day 1. Note how

Stranger Things and BG3 score positive overall (attitude_avg 0.36, 0.32) but still register cash_grab and pandering tags from a minority of voices: that's the kind of nuance a horizontal social-listening tool's average-sentiment scalar erases.

What this would look like for a real UA packet. Same classifier, same three forum sources, same six narrative tags — but pointed at the new harvest filter. Stage 4 of the pipeline (harvest_forum_presence.py) takes search terms; swap "Goblin Slayer" for "Witch UA" or "One D&D Warlock playtest" and you get the same six-column breakdown 24-48 hours after the packet drops, instead of waiting six weeks for the official survey.

The Homebrew Gap Identifier

The question. Which IPs is the D&D community building unofficial 5e content for, despite WoTC having shipped no licensed crossover? The gap between community demand and official supply is the lead.

What we ran. Cross-platform homebrew query against `gold_data.homebrew_combined_proxy`. Every IP in the master list is scored against three independent platforms — D&D Beyond Homebrew (the official creator workshop), GMBinder, and Homebrewery — with the AI Bouncer disambiguating false-positive name collisions on each side.

```
SELECT ip_name, homebrew_combined_score, homebrew_combined_status,
       ddb_total_items, ddb_top_item_name, ddb_top_item_section, ddb_top_item_adds,
       gmbinder_confirmed, homebrewery_confirmed, ua_homebrew_mention_count,
       ua_total_upvotes FROM gold_data.homebrew_combined_proxy WHERE homebrew_combined_status
LIKE 'sufficient%' ORDER BY homebrew_combined_score DESC LIMIT 10;
```

What came back. 60 of 142 candidate IPs cleared the "sufficient" homebrew bar. 5 cleared all three platforms simultaneously. None of the top-10 listed below has a licensed official D&D crossover — every column is unofficial community labor:

IP	Combined score	Streams	DDB items	Top DDB item / adds	GMBinder	Homebrewery
Bloodborne	0.98	2 of 3	71	Bloodborne Hunter / 213	7	2
Hollow Knight	0.94	2 of 3	50	Hollow Knight Vessel / 269	8	1
Elden Ring	0.93	3 of 3	29	Margit the Fell Omen / 38	10	0
Final Fantasy XIV	0.96	1 of 3	0	—	8	1
Demon Slayer	0.81	2 of 3	55	Blood Hunter / Demon Slayer / 718	4	0
One Piece	0.89	2 of 3	28	Cyborg (One Piece) / 75	7	3
Berserk	0.74	3 of 3	25	Berserker Redux / 227	2	0
Mistborn	0.73	2 of 3	44	Mistborn / 117	3	0
Jujutsu Kaisen	0.91	2 of 3	0	—	5	1
Discworld	0.81	1 of 3	0	—	6	0

Source: `gold_data.homebrew_combined_proxy`, snapshot 2026-04-30. "Adds" = the public DDB-Homebrew add count for the top-ranked unofficial item per IP.

What it tells us. Two patterns. First, the all-three-streams set (Berserk, Elden Ring, plus three others not shown) is the highest-conviction gap — the community is building this content on three separate platforms, simultaneously. That's the canonical "we want this and you're not shipping it" signal. Second, the single-item adds count is the diagnostic: "Blood Hunter, Order of the Demon Slayer" has been added 718 times by D&D Beyond users for a campaign — that's not noise, that's a community vote. The Berserker Redux subclass at 227 adds tells WoTC the rage-mechanic remix specifically, not just "Berserk" generically, is the rules-design vector.

What this enables. Pre-licensing pitch: "the Berserk fanbase is already running 25 unofficial homebrew artifacts on D&D Beyond, including a 227-add Berserker subclass — your audience-conversion cost is near zero because they've already converted themselves." Plus the same query, run quarterly, surfaces the velocity changes: which IPs jumped from "sufficient_one" to "sufficient_all_three" in the last 90 days. That delta is the early-warning track for the next licensing window.

The Mechanic Demand Detector

The question. Three independent streams measure each D&D class: search interest (curiosity), community discussion (community), and creator output (creator). When these three diverge for the same class, what does each pattern of disagreement mean for design priorities?

What we ran. Pulled the `composite_concept_index` for D&D's 12 evergreen classes. Each stream contributes a normalized 0-1 score; we additionally surface `curiosity_momentum` and `community_momentum` — the directional 7-day deltas. The reading column is the divergence interpretation derived from the score-spread.

```
SELECT concept_name, curiosity_score, community_score, creator_score,
demand_score, ROUND(community_score - curiosity_score, 2) AS comm_minus_search,
streams_present FROM gold_data.composite_concept_index WHERE category = 'Class' AND
is_active = TRUE AND streams_present >= 3 ORDER BY community_score DESC NULLS LAST;
```

What came back. Twelve classes, three stream scores, plus the divergence reading. The fourth column (Comm – Search) is the diagnostic — positive means community advocates more than the public searches; negative means search-led discovery has outrun community engagement:

Class	Search	Community	Creator	Comm – Search	Reading
Paladin	0.97	0.96	0.96	-0.01	Aligned high — saturated
Fighter	0.98	0.94	0.96	-0.04	Search-led, broad appeal
Warlock	0.71	0.95	0.99	+0.24	Community + creator advocacy
Sorcerer	0.68	0.93	—	+0.25	Community-led — under-marketed
Barbarian	0.93	0.93	0.96	0.00	Aligned
Rogue	0.89	0.90	0.96	+0.01	Aligned
Ranger	0.96	0.85	0.99	-0.11	Search-led discovery
Bard	0.75	0.87	0.99	+0.13	Community + creator
Wizard	0.85	0.85	0.96	0.00	Aligned
Cleric	0.95	0.72	0.96	-0.23	Search-led — newcomer entry
Monk	0.78	0.98	0.99	+0.20	Community + creator beloved
Artificer	0.72	0.83	—	+0.11	Community-led, no creator

Source: `gold_data.composite_concept_index`, snapshot 2026-04-30, ruleset = Core Evergreen except Artificer (2014). "—" = creator stream returned 0 for that class, treated as missing rather than zero.

What it tells us. Three actionable patterns. (1) The community-led classes — Sorcerer (+0.25), Warlock (+0.24), Monk (+0.20) — are talked about more than they're searched for. That asymmetry says the engaged players advocate for these classes, but the public discovery layer is weaker. The implication for design is the opposite of "buff Sorcerer mechanically" — the lever is entry-point marketing or a tentpole-product Sorcerer hook, because the mechanical fanbase is already there. (2) The search-led classes — Cleric (-0.23), Ranger (-0.11) — are what newcomers find when they Google "D&D classes," but the existing community has moved on. That argues for content refresh: keep Cleric/Ranger relevant for the players who already chose them. (3) Artificer reads community 0.83 / creator 0 — the creator stream is empty because no Artificer-specific Itch.io / DMs Guild creator content exists at scale. That's a creator-economy gap, not a player-demand gap, and it's a different lever.

What this enables. A quarterly design read: "this class's three streams diverged by more than 0.20 — that's where to look." The same query, run against subclasses, surfaces the same pattern at finer resolution. Hexblade (community 0.69 / curiosity 0.24, +0.45 spread, 3 streams) is a hidden Hexblade-class-tier signal: established community but no broad search interest — exactly the kind of signal a horizontal social-listening tool would miss because Hexblade isn't trending in the volume sense. Trusight surfaces it because the divergence between streams, not the volume in any one stream, is the lead.

PERSONA A — TAKEAWAY

Three queries, three real findings.

Each of the three Persona-A use cases is a live SQL query against gold_data — not a slide deck description. The classifier produces six narrative tags per forum thread (A1). Sixty IPs cleared the cross-platform homebrew bar with five clearing all three streams simultaneously (A2). Twelve classes scored across three streams reveal three separable divergence patterns, each with a distinct design action (A3).

Most importantly: every claim above traces to a row in a gold-tier view, with snapshot date and source column. That's the data trail Trusight ships — the structured signal that turns design judgment from "my gut says" into "the row in the table says."

Next tranche: Persona B (Marketers) — proof-of-use for OGL-Tier Backlash Detector, Creator ROI Attribution, Audience Segmentation Atlas.