

Business Plan

counterview

See who confirms a claim, who disputes it, and what the debate leaves out.

No verdict — sources shown, you decide.

A Computer Vision Center / Universitat Autònoma de Barcelona spin-off (in incorporation)

Arka Ujjal Dey (Founder & Project Lead) · Ernest Valveny · Dimosthenis Karatzas · Artemis Llabrés

audey@cvc.uab.cat · ernest@cvc.uab.cat

Version 30 · revised 25 October 2026 after feedback round 1 and an expert review (sources, consistency, operating break-even) · Confidential (contains internal costs and price hypotheses)

TABLE OF CONTENTS

0. Executive Summary.....	3
1. The Market.....	5
1.1 The problem and the opportunity.....	5
1.2 Market size and target markets.....	5
1.3 Single solution and proof of concept.....	6
1.4 Competitive structure and prices.....	6
1.5 Clientele structure.....	7
2. Marketing / Merchandising Plan.....	8
2.1 Products and services.....	8
2.2 Opportunities (next two years).....	9
2.3 Price.....	9
2.4 Sales process, international plan and monetization.....	9
3. Development Plan.....	11
3.1 Technical challenge and state of the art.....	11
3.2 Current status.....	11
3.3 Future developments (5 years).....	11
3.4 IP strategy.....	12
4. Promoter Team and Incorporation Plan.....	13
4.1 Promoters and key team.....	13
4.2 Shareholding structure.....	13
4.3 Organisation chart.....	13
4.4 Positions and profiles.....	14
4.5 Incorporation plan (headcount).....	14
4.6 Remuneration system.....	14
5. SWOT Analysis and Contingency Plans.....	15
5.1 PESTEL (external factors).....	15
5.2 SWOT matrix and strategies.....	15
5.3 Contingency plans.....	15
6. Operational Plan (Implementation).....	16
6.1 Challenges.....	16
6.2 Schedule.....	16
6.3 Collaborating entities.....	16
7. Economic-Financial Plan.....	17
7.1 General considerations.....	17
7.2 Investment estimation.....	17
7.3 Revenue estimate.....	17
7.4 Staff and other expenses.....	17
7.5 Profit and loss.....	18
7.6 Balance sheet.....	19
7.7 Cash flow.....	19
7.8 Profitability: VAN (NPV) and TIR (IRR).....	19
7.9 Financing plan.....	20
7.10 Scenarios.....	20
7.11 Investor metrics and use of funds.....	21
Annex I – Abbreviated CVs of key people.....	21
Annex II – Business Model Canvas.....	22
Annex X – Other relevant information.....	23
Immediate action plan: the next 12 weeks.....	24
Targets for week 12.....	24
Phase 1 — Launch pilots (weeks 1–3, 12 Oct – 1 Nov).....	24
Phase 2 — Learn from use (weeks 4–6, 2 – 22 Nov).....	24
Phase 3 — Convert and design (weeks 7–9, 23 Nov – 13 Dec).....	24
Phase 4 — Decide and prepare 2027 (weeks 10–12, 14 Dec – 3 Jan).....	25

0. Executive Summary

At a glance

The ask. €300k from aligned angels in 2027 (convertible, €2.0–2.5M cap), which unlocks €485k of ENISA loans and ACCIÓ and NEOTEC grants. No further round is planned.

Proof so far. A working product (v7) used on real cases; cost per check measured from our spend ledger (€0.09 for a typed claim); 93–96% agreement on typed claims (provisional, not yet against human experts); betevé and 3Cat have seen the product. No pilot or paying customer yet.

Next milestones. First paid pilots (November 2026); incorporation (Q1 2027); June 2027 go / no-go: ≥5 paying customers or €60k contracted ARR, ≥50% pilot retention, ≥80% agreement with a human-expert reference.

The numbers. €1.5M revenue and €1.6M ARR in 2031; gross margin 65–77%; EBITDA positive in 2029 with grant income, 2030 before it; lowest year-end cash €304k.

Main risks. The founder sells alone until the first business developer (2028); willingness to pay is untested; video-claim accuracy is below target; the product depends on search-engine terms that are under legal review.

Company introduction

counterview shows who confirms a claim, who disputes it, and what the debate leaves out. Paste a claim, a social post with its photo, or a video link: counterview lists the checkable claims (in a video the user ticks which to check), searches the open web and your own offline libraries, and labels every source by what it does with the claim — supports, contradicts, related or not relevant, with finer tags (core only, echo, denied, different number) — with the quote and date it was judged on, who made the call, and the reason: which part of the claim the source confirms or disputes (who, what, where, when, how many). **It never gives a verdict: sources shown, you decide.**

counterview is a spin-off in incorporation from the Computer Vision Center (CVC) and the Universitat Autònoma de Barcelona (UAB), built on the team's published research on retrieval-augmented verification ([arXiv:2504.10166](https://arxiv.org/abs/2504.10166)), developed within the Spanish-government REMISS project. Its promoters are the four co-inventors: **Arka Ujjal Dey** (Founder & Project Lead), **Ernest Valveny** and **Dimosthenis Karatzas** (UAB professors, CVC researchers) and **Artemis Llabrés** (CVC technical staff). Today Arka and Ernest lead the work; Dimosthenis and Artemis advise and will choose their roles later. Hiring follows revenue: one engineer from incorporation and a second in 2028 for the on-premise version (grant-funded R&D); a business developer once recurring revenue reaches €75k; a Head of Commercial (salary, commission and 4–6% options, not a co-founder block) once it reaches €500k. The AI accelerator helps recruit them. Two members of the accelerator team already work with us, Marc Rubí (engineering) and Shreyanshi Tiwari (business development). Their long-term roles, as early employees or co-founders, will be agreed together at the end of the programme, with any equity coming from the 15% option pool with four-year vesting.

Business opportunity

Problem. A claim goes viral and teams have minutes, not days. Search engines return links; fact-check sites cover a fraction of claims, days later; media-monitoring suites count mentions; threat-intelligence tools give risk scores without showing the evidence. Nobody shows the full spread of coverage — who confirms, who disputes, where a photo first appeared, what a 20-minute video actually claims — fast and without taking a side.

Solution and customers. One workspace, sold in three steps: (1) **newsrooms and fact-checkers** — "both sides in minutes, with the quotes — you keep the verdict"; (2) **PR and communications teams** — "know who is saying what about you, and who disputes it, before you respond"; (3) **public-integrity and police OSINT units** — "runs on your server, nothing leaves the building; trace where a photo or video first appeared". An **API** follows from 2028. Markets: Spain, the rest of the EU and the UK. counterview is an evidence and investigation tool, not a monitoring suite: it is priced on the analyst hours it replaces, not against social-listening subscriptions.

Current status

A working product (v7, October 2026) with accounts, admin panel, exports, an evidence pack, published fact-checks of the same claim (Google Fact Check Tools) and offline topic packs. Real checks: a 20-minute political video → **95 claims, 934 source items, 12 blind spots**; a viral Hurricane Helene photo → **2 outlets confirming vs 14 disputing**, earliest report traced to 28 September 2024; a photo shared as a Delhi protest → **seven sources placing it in Mumbai and Thane**, each naming the place. Accuracy on 373 sources for 20 typed claims: the direction we give each source

matches our reference **93–96%** of the time and support and contradiction are swapped in 0.5–1.9% (provisional: the reference is not yet checked by a human expert). Video claims are the weak point — **the wrong side 26–33% of the time** on a small single-annotator set — and are being fixed with user corrections and a larger human-labelled set. Cost per check at vendors' list prices: **€0.09 for a typed claim, €0.17 for an image post, €0.07–0.11 per video claim checked** and €0.01 offline; web search is ≈92% of it. **No pilot has started yet (first cohort November 2026)**; betevé and TV3/CCMA have seen the product. No paying customer yet.

Key figures (base case) and funding




Figure 1. Revenue, EBITDA and headcount, base case (€ thousands). Full model in Section 7.

Revenue grows from €93k (2027) to €1,497k (2031), reaching €1.6M ARR with 80 newsrooms, 48 PR agencies and 10 public-integrity units. Every customer starts with a paid pilot: one month for €500 for newsrooms and fact-checkers, two months for €2,500 for PR agencies and campaign teams. Plans then cost €475 a month for newsrooms and €1,200 a month plus €350 per active crisis client for PR agencies; public-integrity bodies pay from €15k a year in the cloud or €35k on-premise. Annual prepayment earns 8% off. Plans are sold in check credits, so usage stays ≤15% of every plan's price; gross margin is 65–77% (SaaS basis, including hosting, customer success and pilot delivery). The company is built to be self-sustaining: each hire needs a recurring-revenue trigger and at least three months of costs in the bank afterwards, and the founder's salary rises only once recurring revenue passes €300k. EBITDA turns positive in 2029 (€51k, including €110k of NEOTEC grant income; before grants, operating break-even is 2030), stays about flat in 2030 while the Head of Commercial and a third engineer join, and reaches €196k (13% of revenue) in 2031. Funding: €300k from aligned angels in 2027 (convertible loan, €2.0–2.5M cap, information rights only) unlocks €485k of non-dilutive finance — ENISA loans (€175k), an ACCIÓ Startup Capital grant (≈€90k) and a CDTI NEOTEC grant (€220k). No further equity is needed: the lowest year-end cash is €304k, and the downside case (half the customers, a year later) keeps at least €166k because hiring waits for revenue. Founders keep ≈69% (≈82% with the option pool). Newsrooms are our first paying segment and our design partners; PR and crisis agencies and public-integrity units are the growth engines.



Figure 2. Key milestones, Q4 2026 – 2031 (the next 12 weeks are planned in the final section).

1. The Market

1.1 The problem and the opportunity

Viral claims arrive as text, photos and video, and spread far faster than corrections: the truth takes about six times as long as falsehood to reach 1,500 people ([Vosoughi, Roy & Aral, Science, 2018](#)). Three kinds of team must react within minutes: **newsrooms and fact-checkers** (is this true, who has already rebutted it?), **PR and communications teams** (is this attack on our client spreading, and who disputes it?) and **public-integrity and police OSINT units** (is this photo recycled from another event, where did it first appear?). Their tools do not answer these questions: search engines give links, fact-checks arrive days later, monitoring suites count mentions, and AI assistants give one confident answer that cannot be audited.

Future trends

- **AI-generated and recycled media** multiply the volume of photos and videos to check; provenance (where an image first appeared) becomes routine work.
- **Fact-checking budgets are under pressure**: only 22.6% of fact-checking organisations describe their finances as sustainable and 76% as vulnerable or in crisis ([IFCN, State of the Fact-Checkers 2025](#)) — tools that save checking time are needed, at newsroom prices.
- **Regulation** — the EU Digital Services Act, the European Media Freedom Act and the AI Act — raises expectations of traceable, auditable decisions about content.
- **Communications teams already pay for monitoring at scale** (a median Meltwater contract is about €22,400 a year; [Vendr data](#)), but monitoring counts mentions. Investigating a claim — transcribing, cross-referencing outlets, finding the quotes — is still done by hand: a fact-check takes from 15–30 minutes for the simplest claim to a full day for a typical one, and two or more days for a complicated one ([Hassan et al., 2015, co-written with PolitiFact's founder](#)), and tracing where an image came from takes about 50 minutes ([Caulfield, Nieman Lab, 2026](#)). That labour, not the monitoring budget, is what counterview replaces.
- **Election cycles** ([Spain 2027](#), European Parliament 2029, UK general election by 2029) create recurring peaks of demand.

Barriers to entry and how we overcome them

Barrier	How counterview overcomes it
Trust in AI output	No verdicts; every label shows its quote, date, who made the call and the judge's reason (what the source confirms or disputes: who, what, where, when, how many); typed claims match our reference 93–96% (provisional; human-expert check pending); published method.
Workflow fit	Paid pilots (one month for newsrooms, two for PR) measure activation, usefulness and retention before a full plan; exports to newsroom reports, CSV and publishable pages.
Language	Spanish and Catalan interface before broadcaster use; multilingual transcription already tested in English and Spanish.
Security and confidentiality (police)	Offline topic packs and local models; chain of custody, two-factor login and a stronger local judge before any police sale.
Cost of AI calls	Cost per check measured at list prices (web search ≈92% of it): €0.09 per typed claim, €0.17 per image post, €0.07–0.11 per video claim checked, €0.01 offline. Plans are sold in check credits, so usage cost stays ≤15% of every plan's price at full use; gross margin ≥70% with hosted AI models.

1.2 Market size and target markets

Total addressable market (TAM): media monitoring tools, €4.7 billion in 2024 (USD 5.4 billion at [1 EUR = 1.15 USD; market study by Viromii for CVC, March 2026](#), p. 26), growing 15.5% a year ([Grand View Research](#)) — shown as market context only: we do not price as a monitoring replacement; our prices are anchored on the analyst time a check saves (§2.3). AI content verification tools, €3.3 billion (USD 3.8 billion) in 2024, growing 21% a year ([Grand View Research](#)), are an adjacent category shown as growth context only: the study warns that the two overlap, so they are not added together. Serviceable addressable market (SAM): a bottom-up count of buyers in Spain, the rest of the EU and the UK, at our prices:

Segment (EU + UK)	Buyers	Annual value	SAM
Newsroom verification teams and fact-checkers (broadcasters, publishers, fact-checking organisations)	~2,000 teams	€5.7k	€11M
PR agencies, in-house communications teams and public-campaign managers with issues and crisis work	~5,000	€15.5k	€77M
Public-integrity, electoral and police OSINT units	~300	€30k	€9M
AI product teams using the API (from 2028)	~3,000	€5k	€15M
Total SAM			≈ €113M

Buyer counts are team estimates, to be checked during the pilots, anchored on public counts: [115 member organisations in the EBU \(2026\)](#), 437 active fact-checking projects worldwide ([Duke Reporters' Lab, June 2026](#)), more than 3,000 PR firms in the [national associations of ICCO](#) and 74,460 public-relations and communication enterprises in the EU-27 ([Eurostat](#), 2020, cited in the market study). Prices are the averages in \$2.3.

Obtainable market (SOM). The base case reaches 80 newsroom teams, 48 PR and communications agencies and 10 public-integrity units by 2031, plus API usage: €1.6M ARR, about 1.4% of the SAM. Newsrooms are the first paying segment and our design partners; PR and crisis agencies and public-integrity units carry the growth. Sales start in Spain and Catalonia, extend to the UK from 2027 and the rest of the EU from 2028, and to police and public-integrity units in the EU and UK from 2028.

1.3 Single solution and proof of concept




Figure 3. One engine: four kinds of input, one labelled view of the evidence.

Where a search engine returns links, a fact-checker returns a verdict and a monitoring suite returns mentions, counterview returns **the spread of coverage around each claim**: every source, labelled by what it does with the claim, with the quote, the date and who made the call; the earliest report; how coverage spread over time; and, for debates and interviews, the topics nobody on stage addressed. Each source *supports*, *contradicts* or is *unrelated* to the claim, with finer tags — a supporting source may be direct evidence, an echo of the claim or confirm the *core only*; a contradicting one may be direct evidence, a *denial* or give a *different number* — so partial truths that binary true/false tools miss stay visible.

Real check (October 2026)	Input	Result
"India lost five planes including three Rafales in May 2025"	Text claim	21 outlets — India admitting losses, Pakistan's claims and analysts' different figures, each with its own quote
Rescue-boat photo shared during Hurricane Helene	Photo + caption	15 independent outlets: 2 confirm, 14 dispute; earliest report 28 Sep 2024
Interview with India's Chief of Defence Staff (3.5 min)	Video	11 claims by speaker, linked to timestamps; 44 outlets
20-minute US political video	Video	95 claims (45 worth checking), 6 topics, 934 source items, 12 blind spots

1.4 Competitive structure and prices

No competitor shows, per claim, who confirms and who disputes it, with the quote behind each label, across text, images and video, and without a verdict. Published prices (or price evidence) show what buyers already pay for adjacent tools. We do not price counterview as a replacement for a monitoring suite: it does not count mentions or clip coverage, so its price is anchored on analyst time saved (\$2.3).

Player	What it does	Price evidence	Gap vs counterview
Meltwater	Media monitoring and social listening	Quote only; €4.7–49.6k/yr across 126 purchases, median ≈€22.4k (\$25,800; Vendr, Feb 2026)	Counts mentions; no stance per source
Cision	PR monitoring and distribution	SMB ≈ €9.2k/yr; enterprise ≈ €76k/yr (SpendHound, Oct 2026)	Same
Brandwatch · Talkwalker	Social listening	No public price; quote only (Toolradar)	Sentiment, not evidence
NewsWhip	Predictive media monitoring	Per user, quoted per contract (Capterra)	Predicts spread; no claim labels
Factiverse	AI fact-checking for newsrooms	On request (iTechGuides)	Verdict-oriented; text-first
InVID-WeVerify plugin	Free verification toolbox (video keyframes, reverse image search, metadata) used by verification desks	Free	Tools for a person to investigate; no source-by-source labels or evidence brief
Google Backstory	AI tool that traces an image's history for fact-checkers (2026)	Not public (Nieman Lab, 2026)	Images only; no stance per source; no text or video claims
AI answer engines (Perplexity, ChatGPT search)	Answers with citations	Per-user subscriptions	One synthesised answer; no per-source stance, quote or audit trail
Self-serve AI fact-check tools	Claim checks for individuals	€13–39/month (ViralBuster)	Single answers; no team workflow, images or video
Claim-check and image APIs	Building blocks	≈ €7 per 1,000 claims (Apify); TinEye €0.01–0.035 per image search (TinEye)	Components, not a workspace
Recorded Future · Maltego	Threat intelligence and OSINT	Recorded Future from ≈€23k/yr, median ≈€63k (Vendr); Maltego pricing on request	Investigation graphs; no claim-level coverage
Blackbird.AI · Cyabra	Narrative threat intelligence	Enterprise contracts	Risk scores and bots, not neutral evidence

Prices are third-party estimates from public pages, retrieved October 2026; vendors negotiate.

Why won't Google or an AI answer engine do this? Free tools help a person investigate (InVID-WeVerify, reverse image search) or give one synthesised answer (AI answer engines); Google's Backstory traces an image's history. None labels every source by what it does with the claim, with the quote, date and reason, as an evidence trail a desk can publish or file. Our defence is that workflow, the label scheme and expert-labelled data behind the stance judge, offline and on-premise deployment for public bodies (nothing leaves the building) and being an EU provider for EU public buyers. If a large platform ships claim-level evidence views, we lean further into on-premise, topic packs and public-integrity work.

1.5 Clientele structure

Segment	Buyer	Pain	What we sell	Sales cycle
1. Newsrooms and fact-checkers (first paying segment; design partners)	Head of verification, standards editor; IFCN and EFCN members	Slow sourcing; one-sided checks; viral images and video	Team plan (seats + monthly checks); event packages for elections	1–3 months; small budgets, grants (Google News Initiative, EU EDMO / EMIF)
2. PR, communications and campaigns (growth engine)	Agency account director; in-house head of communications; public-campaign manager	Is this attack spreading? Respond or stay quiet? Who carries it?	Agency workspace (3 seats) plus a per-crisis add-on for each client under attack; claim tracking and alerts; evidence briefs; outlet and framing maps	Weeks; entry price below most approval thresholds
3. Public-integrity and police OSINT units (growth engine; cloud now, on-premise from 2028)	Head of unit; procurement	Viral misinformation, recycled footage, evidence trail	Cloud licence for electoral and integrity bodies; on-premise or offline licence for police and defence (local models, case export, training)	6–18 months; tenders, security review
4. AI product teams (API, from 2028)	CTO / product lead	Evidence behind AI-written answers	Pay-per-check API	Self-serve

2. Marketing / Merchandising Plan

2.1 Products and services

One workspace, packaged for three kinds of team, plus an API later. Every package uses the same engine and the same labels.

Package	For whom	What it includes	Status
Newsroom	Verification desks, fact-checkers, broadcasters	Text, image-post and video checks (the user ticks which claims of a video to check); every source labelled with its quote and the reason; earliest report; published fact-checks of the same claim; framing map for long videos; newsroom report, CSV, evidence pack and publishable-page exports; election and event packages with prepared topic packs. 600 check credits a month (1 per claim, 2 per image post); checks answered from the offline library or a topic pack use no credit	Working (v7); Spanish/Catalan interface and credit metering Q4 2026
PR & crisis communications	Agencies and in-house teams	Agency workspace: 3 seats and 1,000 check credits a month. Per-crisis add-on for each client under attack: claim tracking (20 claims re-run every 6 hours, alerts when new outlets dispute or repeat a claim) and an evidence brief — every outlet that repeated the claim, who denied it, the qualifying quotes; outlet table, framing map and client-ready page. No social listening or mention counts	Claim tracking and alerts: Q1–Q2 2027
Public-integrity / OSINT	Police cyber units, electoral and integrity bodies	Cloud licence for electoral and integrity bodies (2,500 live credits a month, pooled annually), or offline / on-premise install for police and defence; local models; offline topic packs; image history (first-seen copy, same photo / same scene); case export with chain of custody; two-factor login; training	2028 (on-premise needs chain of custody, local judge ≥75%, legal review)
API	AI product teams, platforms	Checks as a service: claims and labelled sources as JSON, paid per credit; standard or economy search; hard limits per check (searches, sources, time, cost)	2028

The screenshot displays the counterview interface for a workspace titled "Detention of Cockroach Movement Leaders in Delhi". The interface is divided into several sections:

- Top Bar:** Includes the workspace name, a "Live web off" indicator, and buttons for "Services", "New check", "Share", and "PR".
- Left Sidebar:** Contains a "Checks" section with a search bar and a list of checks, and a "Packs" section with a "prep topic" button and a note about prepared topics.
- Main Panel:** Displays a list of claims under the "Claims" tab. Each claim entry includes a timestamp, a title, a progress bar, and details about support and contradicting evidence. For example, one claim states "130 million voters across India were removed from the voter list in this update." with 37 outlets, 36 independent, and 47 support.
- Right Panel:** Contains tabs for "Claim", "Sources", "Searches", "Media", "Narrative", "Engine", and "PR". The "Claim" tab is active, showing details for a selected claim, including a "Support: Core only (preliminary)" section, a "Contradict: Evidence (preliminary)" section, and a "Contradict: Evidence (preliminary)" section.

Figure 4. counterview on a 43-minute France 24 video about India's youth-led protests: 367 claims, filtered to the 36 with sources on both sides; the panel shows, for the selected claim, each source's quote, what it confirms or differs on (who, what, where, when, how) and the reason for its label.

2.2 Opportunities (next two years)

Opportunity	Status	Proposal	Expected €	Timing
Paid design-partner pilot cohort (newsrooms, fact-checkers, PR teams)	Not started yet; first cohort planned for November 2026 (newsrooms: 1 month, €500; PR: 2 months, €2,500)	Pilot → paid plan	€500–2,500 per pilot; then €5.4–6k/yr (newsroom) or ≈€15k/yr (PR)	Conversions from Q1 2027
betevé and TV3/CCMA	Saw the product in 2026; meetings October 2026	Join the pilot → election package for May 2027	€12–14.5k package (price hypothesis)	Q4 2026 – Q2 2027
PR agencies and public-campaign managers (Spain, UK)	Targets: PR agencies and managers of public and political campaigns; first contacts through the accelerator and mentors	Discovery calls → paid crisis pilot	€2,500 pilot, then €1,200/month + €350 per active crisis client	From Q1 2027
Fact-checkers (Verificat, Maldita.es, Newtral, EFE Verifica, Full Fact)	Design partners and pilot users	Discounted team plan; research collaboration	€3–6k/yr	2027
First public-integrity or police unit (EU/UK)	Exploratory conversations only until legal review	Cloud or offline pilot → licence	€15–25k/yr cloud; €35–50k/yr + €10k on-premise	2028

Before incorporation, any paid work is contracted by CVC as a research contract.

2.3 Price

Prices are anchored on the analyst time a check replaces, not on monitoring-suite prices: counterview is an investigation and evidence tool, not a social-listening dashboard, and a buyer who expects mention counts would churn. Every customer starts with a paid design-partner pilot — €500 for one month (newsrooms and fact-checkers), €2,500 for two months (PR agencies and campaign teams), public-integrity pilots agreed case by case — long enough to judge usefulness and short enough to decide. [Plans then cost €475 a month for newsrooms \(600 credits\) and €1,200 a month for PR agencies \(3 seats, 1,000 credits\) plus €350 a month for each active crisis client \(agencies can bill it to the client's crisis retainer\); public integrity from €15k a year in the cloud \(2,500 credits a month, pooled annually\) or €35–50k a year on-premise plus €10k installation. Annual prepayment earns 8% off. Public buyers get one contract per year below the €15k contrato-menor cap \(excl. VAT\): an election or event package includes the team plan for the campaign months, and plan months already paid that year are credited against it. Prices appear in the client deck and in meetings, not on the website, which only collects pilot requests; our costs are shown only to investors. Every plan includes a monthly allowance of check credits, the unit that drives our cost: one credit per claim checked against the live web \(a typed claim, or a claim the user ticks in a video\), two per image post. Transcribing a video and listing its claims is included; checks answered from the offline library or a prepared topic pack use no credit. Usage above the allowance triggers a soft alert, then extra credits at €0.40–0.50, so clients get no surprise bills and heavy use cannot erode the margin. Custom offline topic packs \(a prepared evidence library for an event, election or investigation\) are included in election packages and public-integrity licences; for other customers they are priced at about half the researcher time they save — roughly €300–600 a pack against 2–5 days of a researcher assembling sources \(≈€570–1,430 at ≈€36 an hour\). They are upside, not part of the revenue plan.](#)

Value anchor. A fact-check takes from 15–30 minutes for the simplest claim to a full day for a typical one, and two or more days for a complicated one ([Hassan et al., 2015, co-written with PolitiFact's founder](#)); later reviews repeat the range ([Fan et al., EMNLP 2020](#); [Thorne & Vlachos, COLING 2018](#)). Tracing an image's history by hand takes about 50 minutes ([Caulfield, Nieman Lab, 2026](#)). An hour of a high-skill employee in Catalonia costs the employer about €36: €27.37 gross pay per hour ([INE, Encuesta Anual de Estructura Salarial 2024](#)) plus 30.65% employer social security ([BOE, Orden PJC/297/2026](#)). At the simplest-claim minimum of 15 minutes, one typed claim therefore costs about €9 of analyst time, an image about €30, and the 10 claims we pre-select in a 20-minute video about €89. counterview's usage cost is €0.09 per typed claim, €0.17 per image post and €0.07–0.11 per video claim ([our spend ledger at list prices](#)). Even if counterview cut checking time by only 20% (the gain measured for short evidence briefs by Fan et al.), a newsroom using its 600 monthly credits on typed claims would save about €1,000 of analyst time a month against a €475 plan. If an analyst instead reviews counterview's brief in about 5 minutes rather than checking for 15 (our working assumption, which the pilots will measure), checking time falls by 67% and the cost of checking, including our fee, by about 58% for a newsroom (600 claims, €475 plan) and 53% for a PR agency (1,000 claims, €1,200 plan). INE publishes no wage for journalists or PR specialists separately, and its 2022 pay averages for media and for advertising and consultancy differ by about 12% ([INE sector pay 2022](#)), so one rate is used for both; at €30–45 an hour the reductions stay at 56–60% and 51–56%; automated image-history tools cut 50 minutes to 3 (Caulfield). The calculation is in [references/R00 cost per check calculation.xlsx](#). The pilots measure each team's own time per check (“How long would this check have taken your team manually?”).

Package	Price (excl. VAT)	Included usage	Competitor reference	Gross margin at full use*
Newsroom team plan	€475/month list; 3–5 seats; 1-month pilot €500; 8% off for annual prepayment	600 credits a month; soft alert, then €0.50 per extra credit	Above self-serve AI tools (€13–39/month); review in 5 min vs 15 by hand: 67% less time, ≈58% lower cost of checking (≈€36 an hour, INE 2024 + 2026 social security)	81% (usage 15%)
Election / event package	€12–14.5k per campaign; includes the team plan for the campaign months	Up to 10,000 credits + prepared topic packs	Under the €15k contrato-menor cap : one contract per public client per year	85% (usage 11%)
PR & crisis communications	€1,200/month per agency (3 seats) + €350/month per active crisis client; 2-month pilot €2,500; 8% off for annual prepayment	1,000 credits; claim tracking for crisis clients (20 claims every 6 hours); soft alert, then €0.40 per extra credit	Not priced against monitoring suites (Meltwater median ≈€22k/yr): review in 5 min vs 15 by hand: 67% less time, ≈53% lower cost of checking (≈€36 an hour, INE 2024 + 2026 social security)	84% (usage 12%)
Public-integrity / OSINT licence	Cloud: €15–25k/yr (electoral and integrity bodies). On-premise: €35–50k/yr + €10k set-up; hardware bought by the client. Blended ≈€30k/yr	Cloud: 2,500 live credits a month, pooled over the year; offline searches unlimited. On-premise: unlimited offline checks on the client's server; monthly topic-pack refresh	OSINT platforms such as Maltego (enterprise seat); Recorded Future from ≈€23k/yr (Vendr)	86% (pack refresh ≈10%) before deployment staff
API	€0.50 per credit (standard search); €0.20 economy (Brave only)	Pay per use; hard limits per check	Claim APIs ≈€7/1,000; TinEye €0.01–0.035/image	72% standard · 76% economy
Custom offline topic pack	≈€300–600 per pack (about half the researcher time it saves); included in election packages and public licences	Prepared evidence library for an event, election or investigation	Replaces 2–5 days of a researcher assembling sources (≈€570–1,430)	≈65% (half a day of our time + ≈€2 compute)

*At full use of the allowance, list API prices and a hosted AI judge (no on-premise model), after a 2% royalty to UAB/CVC and 2% payment fees, before fixed production hosting and customer-success staff — which the P&L also counts in cost of revenue (§7: 65–77% blended gross margin). Typical use is below the allowance, so realised margins are higher. Sales commissions (4%) are a sales and marketing expense.

Unit economics: what a check costs and why we sell credits

Cost per check. Measured on the current product (v7 with the claim picker) [from our spend ledger](#), 12–18 October 2026, and re-priced at vendors' list prices ([1 EUR = 1.15 USD](#)): [Google grounded search €30 per 1,000 prompts \(Gemini 2.5\)](#), [Brave Search €4.35 per 1,000](#), [Google Vision €3.04 per 1,000](#), the hosted AI judge as billed. Web search is ≈92% of the cost and the AI judge ≈8%; transcription, embeddings and image matching run on our own GPU.

Check	Live web search (list prices)	Offline (library or topic pack)	Credits	Basis
Typed text claim	€0.09 (€0.09–0.13)	€0.009	1	8 checks
Image post with caption	€0.17 (€0.09–0.18)	≈€0.02 (estimate)	2	7 checks
Video: each claim checked	€0.07–0.11	≈€0.01	1	2 videos, 10 and 19 claims checked
Video: transcription and claim listing	€0.004 per minute	€0.004 per minute	Included	Own GPU + extraction model
Planning cost per credit	€0.12 (median + 30% buffer)	—	—	Covers retries, hard claims, extra search rounds

Usage cost of each plan at full use of its allowance:

Plan	Usage at full allowance	Cost at list prices	Share of price	With economy search†
Newsroom plan (€475/month)	600 credits	€72	15%	€24 (5%)
PR workspace (€1,200/month)	1,000 credits + claim tracking (≈€21)	€141	12%	€61 (5%)
Election / event package (€13k)	10,000 credits + topic packs (≈€200)	≈€1,400	11%	≈€600 (5%)
Extra credits	Sold at €0.40–0.50	€0.12	24–30%	€0.04 (8–10%)
API, per credit	€0.50 standard / €0.20 economy	€0.12 / €0.04	24% / 20%	—
For comparison: v28 newsroom design	250 text, 50 image, 120 video minutes at €0.73	€119	24%	—

†Economy search uses Brave only (7 searches and the AI judge: ≈€0.04 a credit); its effect on coverage is not yet measured (a ≈€2 test on 20 claims is scheduled). Our ledger books Google at €0.017 a call, below its list price; we plan on list prices. Two upsides are not in the plan: Google's free tier (1,500 grounded prompts a day) and [Gemini 3 grounding at €12 per 1,000 queries](#). Figures come from one week of development use by one team and will be re-measured in the pilots.

What changed since v28. v28 priced video by the minute because every claim of a video was searched (€0.73 a minute). The claim picker now lets the user tick which claims to check (ten are pre-selected), so cost follows claims checked, not length: a 10-minute video with 19 claims checked cost €1.37 at list prices, against €7.30 at v28's rate. The plan therefore (1) sells check credits instead of video minutes; (2) sizes each plan so that usage is ≤15% of its price at full use; (3) keeps Brave-first economy search and topic packs as cost levers, not plan assumptions; and (4) will enforce credit weights and a cost cap per check (account allowances exist; per-check caps are to build in Q4 2026). The model keeps usage shares above the full-use estimates — newsrooms 15%, PR 15%, public-integrity 10%, API 25% — and adds production hosting, customer-success staff, the royalty and payment fees to cost of revenue. Gross margin then stays at 65–77% in every year of the base case with hosted AI and list search prices; in 2031 it would be 63% if search prices rose 50%, and 61% if every credit cost twice our planning figure (\$7.10).

2.4 Sales process, international plan and monetization

Sales process

- **Pilot first** — a paid design-partner pilot: one month for €500 (newsrooms and fact-checkers) or two months for €2,500 (PR agencies and campaign teams); public-integrity pilots are agreed case by case. Before incorporation, pilots are contracted through CVC as research contracts. Sample checks are pre-loaded and users give one-minute feedback every 5 checks; the first cohort is planned for November 2026 (none has started yet). Go/iterate criteria already instrumented: ≥70% of users activated, ≥60% useful-check rate, ≥50% week-2 retention, usability scores (UMUX-Lite, SUS) ≥68, ≥80% pass rate on pilot tasks.
- **Convert** — two questions after the pilot's checks: “How long would this check have taken your team manually?” and “If this workspace cost €475 a month (newsroom) or €1,200 a month (PR), would that fit your existing team budget, or need senior approval?”; monthly or annual plan at the end.
- **Newsrooms** — direct sales by the founders and business developer; grants (Google News Initiative, EDMO/EMIF) can fund fact-checkers' plans.
- **PR** — agency by agency: a two-month paid pilot, then a €1,200-a-month workspace and a per-crisis add-on the agency bills to its client during a crisis retainer. The pitch: when an attack hits, a verified evidence brief — every outlet that repeated the claim, who denied it and the exact quotes — before the morning press statement.
- **Public integrity** — electoral and integrity bodies on the cloud licence, often below the €15k tender threshold; police and defence units on-premise from 2028, after legal review, through tender or framework agreement.
- **After-sales** — onboarding, (i) help on every control, usage reviews, admin panel for each customer.

International development

Spain and Catalonia first (language, references, broadcaster contacts); the UK from 2027 (large PR market, English-language product); the rest of the EU from 2028 through fact-checking networks and public-integrity bodies. India and other markets are out of scope for now.

Monetization and 5-year sales forecast




Figure 5. Revenue by customer line, base case (€ thousands).

Customers at year-end	2027	2028	2029	2030	2031
Newsroom team plans	8	20	35	55	80
Election / event packages (in the year)	3	2	4	2	3
PR & communications agencies	3	10	20	32	48
Public-integrity / OSINT licences	0	1	3	6	10
ARR at year-end (€k)	85	293	609	1,036	1,606

Revenue (€k)	2027	2028	2029	2030	2031
Newsrooms (avg teams × €5.2k after the prepay discount + packages × €13k)	60	99	196	262	393
PR & communications (avg agencies × €14.3k: €1,200/month after the prepay discount + crisis add-ons)	21	93	215	372	572
Public-integrity / OSINT (avg licences × €30k cloud / on-premise blend + €5k avg set-up)	0	20	70	150	260
API	0	15	50	110	200
Paid pilots (customers who convert)	12	26	39	52	72
Total revenue	93	253	570	946	1,497
Gross margin (cost of revenue incl. hosting and customer success)	77%	74%	65%	70%	69%
EBITDA (€k)	(50)	(19)	50	6	196
EBITDA margin	n/m	-8%	9%	1%	13%

Margin and EBITDA intention. Gross margin stays at 65–77% every year with hosted AI models and list search prices (≥65% is the floor we manage to: credits, pack-first answers and economy search are the levers). EBITDA turns positive in 2029 (2030 before grant income) and reaches 13% of revenue in 2031 (Rule of 40 near 71: 58% growth + 13% margin), with no equity beyond the €300k round.

3. Development Plan

3.1 Technical challenge and state of the art

Checking a claim against the web is a retrieval and labelling problem: find the sources that matter, then say what each source does with the claim. Most tools reduce this to true/false. Real coverage is messier — a source may confirm the event but give a different figure, or simply report what someone said. counterview gives each source a direction — supports, contradicts, related or not relevant — with finer tags (evidence, echo or core only; evidence, denied or different number), records whose words decided it, keeps the quote it was judged on and states the reason: which part of the claim the source confirms or disputes (who, what, where, when, how many, how, why). For images it matches the photo first (same photo, same scene), finds where it first appeared and its original context, then checks the caption against the pages that carry it. The published CRAVE method behind it leads real-world multimodal benchmarks ([arXiv:2504.10166](https://arxiv.org/abs/2504.10166)).

Where we stand (our own numbers, 18 October 2026)

Measure	Result	Comment
Direction of each source (supports / contradicts / unrelated), typed claims	93–96% (373 sources, 20 claims; mean 94% over 8 runs)	Provisional: reference labelled by a separate model and reviewed by us; human-expert check next. Target ≥90%
Support and contradiction swapped, typed claims	0.5–1.9% (mean 1.2%)	Target ≤2%
Video claims: source given the wrong side	26–33% (31 claims, one annotator)	Our weakest point: spoken claims are vaguer than typed ones. Fix: user corrections feed the judge (built); ≥100 human-labelled video claims. Target ≤10%
Finer labels	"Different number" 97–100%; core only, echo, denied 76–81%	Shown in the app as "preliminary"; target ≥90%
Published fact-check of the same claim (Google Fact Check Tools)	Found for 6 of 77 sample claims; 0–1 wrong matches after a same-claim check (5 of 9 before)	New, October 2026; coverage depends on fact-checkers
Earlier 7-label scheme vs our experts	77% (57 labelled sources); videos 74–78%	Experts agree with each other 95%
WiCE (partial support in Wikipedia citations)	62%	Up from 50% in July 2026
Claim extraction	43% of claim cards have at least one unresolved reference ("he", "this deal"); 16% start with one	Measured 12 October; decontextualisation is the first fix
Local model as judge (Qwen3-14B)	54% agreement with the hosted judge (237 judgements, 9 Oct, older scheme)	Must reach ≥75% before offline sales; re-test on the simpler label task
Image provenance (5Pils, 20 images)	Specific location 11, original context 8 (were 8 and 5); publish date 30%	First-seen search to improve
Perspectrum	23%	Not a suitable benchmark — see below

Why we are replacing Perspectrum. Perspectrum consists of abstract opinion statements ("Is nuclear energy ethical?"), not checkable news claims, and our engine is now tuned for news coverage — which explains the drop. Our results on real-world cases have improved over the same period. We will report instead on benchmarks built from real claims with web evidence: **AVERiTeC** (real-world claims), **WiCE**, **QuanTemp** (numerical claims), and our own **200-claim news test set** in Spanish, Catalan and English, labelled by our experts. Perspectrum stays in the research papers, with the explanation. Not yet measured: recall (how many of all relevant sources we find), which newsrooms and police will ask for.

3.2 Current status

- **Product v7 (October 2026)** — text, image-post and video checks with a claim picker (the user ticks which claims of a video to check; the first three start at once). Every source is labelled Support or Contradict with its quote, date, who made the call and the judge's reason (what it confirms or disputes: who, what, where, when, how many); Related sources are shown apart. Published fact-checks of the same claim (Google Fact Check Tools); for images, reverse image search, first-seen date and original context; framing map for long videos; exported reports and an evidence pack (source texts with SHA-256 hashes); accounts with check allowances, admin panel and spend ledger; evidence database of 29,486 passages from 15,854 pages (pgvector, 5 ms per lookup, pages expire after 6 months) and prepared topic packs (e.g. India–Pakistan: 5,939 passages from 308 outlets, built for €1.42).
- **Known limits, stated to clients** — no pilot and no paying customer yet (first cohort planned for November 2026); no verdicts and not court-ready evidence (hashes, but no archived snapshot); no two-factor login; English-only

interface; video claims get the wrong side 26–33% of the time on our small single-annotator set; videos take $\approx 1.5\times$ their length with ten claims checked (target $\leq 0.5\times$); 28–55% of the pages found for videos cannot be read (blocks, paywalls) and are judged on the search summary, flagged (measured 10 October); Google search is rate-limited on some calls; the local judge is weaker than the hosted one; credit weights and per-check cost caps are not yet enforced; the GPU is shared with other research groups (one check failed for lack of GPU memory on 18 October).

3.3 Future developments (5 years)

Period	Development	For
Q4 2026	Spanish and Catalan interface; claim decontextualisation; check credits and per-check cost caps; ≥ 100 human-labelled video claims; Brave-first search test; video time budget (target $\leq 0.5\times$ its length)	Newsrooms, broadcasters
Q1–Q2 2027	Claim tracking and alerts; evidence briefs; client-ready pages; two-factor login and SSO; video wrong-side rate $\leq 10\%$	PR, all
H2 2027	Event mode: prepared topic packs for scheduled debates and elections; API beta	Broadcasters, developers
2028	Chain of custody (archive snapshot + hash per source); local judge $\geq 75\%$ by distillation; offline appliance; security audit and ENS	Police, public integrity
2029	Video copy and crop detection; source-independence (echo) detection; more EU languages	All
2030–2031	Larger offline library; scale infrastructure; API growth	All

Research agenda (papers and grants)

The research team's plan feeds both credibility and grant applications: fine-grained partial stance and "who said it" labels (first papers, with a released seed dataset), claim decontextualisation, source-independence detection, cost-aware search, a small local judge, blind-spot detection, out-of-context images and first-seen provenance, and video copy detection. Target venues include ACL/EMNLP, the FEVER workshop, LREC-COLING, ICWSM and ACM Multimedia.

Technological barriers

- Claims with unresolved references — decontextualisation against speaker and transcript.
- Search cost and dependence on Google — Brave-first search, offline packs and a negotiated quota; the local judge removes the AI-provider dependency ($\approx 10\%$ of cost).
- Long-video processing time — Google search off the critical path, judging as results arrive, a time budget per run, and topic packs.

Participation in R+D+i programmes

Programme	Year	Amount	Role / status
REMISS (Spanish-government research project; 3Cat/CCMA a partner)	Before 2026	Project funding to the research group (not to the company)	Origin of the research
Xarxa Innovation lab	2022	In kind	Market and Patentability Study
Xarxa Innovation lab	2024	In kind	Market Study
AI accelerator programme (4th edition)	2026–2027	In kind	Team building: helps recruit the engineer and business developer once funded; mentoring
ACCIÓ Startup Capital (planned)	June 2027 call	$\approx \text{€}90\text{k}$	Grant, 75% of eligible costs
CDTI NEOTEC (planned)	2027 call	$\text{€}220\text{k}$	2028–2029 project
Torres Quevedo (optional)	2028	Part of a PhD engineer's salary	Upside; not in the model

3.4 IP strategy

Counterview is protected by trade secrets and data rather than by a patent. Rather than relying on public patents that expose algorithmic architecture, counterview is protected by proprietary trade secrets, specialised expert-labelled datasets, custom stance judges and the integration of the multimodal CRAVE pipeline. In AI software, workflow capture, curated evaluation sets and low latency are more defensible than easily bypassed software patents. The research result "CVRAV" was declared to UAB on 7 March 2025 by the four co-inventors (Valveny 25%, Karatzas 25%, Dey 40%, Llabrés 10% — shares of any university income) and is co-owned by UAB and CVC. An early European filing (EP25382241) was rejected because the team's own preprint counted as prior art; future inventions will be filed before any preprint is published. The company will license the software, know-how and expert-labelled data,

proposing equity rather than a revenue royalty — or a 2% royalty that sunsets once cumulative revenue reaches €1M, with a €50–75k buyout option when institutional funding closes. Until UAB/CVC agree, the model conservatively charges the 2% royalty every year.

In practice, the moat is the label scheme and expert-labelled data, the offline evidence library, the workflow built into customers' teams, and speed of execution. Compliance: GDPR impact assessment (content can reveal political opinions); EU AI Act classification; face search off by default and limited to public figures or a legal mandate; terms of use that forbid targeting critics; legal review of search-engine terms (storing and reusing results) and press-content licensing before selling; ENS in 2028, ISO 27001 in 2029.

4. Promoter Team and Incorporation Plan

4.1 Promoters and key team

Person	Role	Background	Dedication
Arka Ujjal Dey	Founder & Project Lead	CVC postdoctoral researcher; first author of the RAV and CRAVE papers; built the engine and product	Full-time from incorporation
Ernest Valveny	Co-founder & Scientific Lead	Professor, UAB; CVC researcher; contact researcher for the invention	Part-time; working on the project now with the Founder & Project Lead
Dimosthenis Karatzas	Co-founder; advisory for now, role to be decided	Professor, UAB; CVC researcher	Part-time; advisory
Artemis Llabrés Brustenga	Co-founder; advisory for now, role to be decided	CVC technical staff	Part-time; advisory
Engineer	Engineering	To be recruited with the AI accelerator once funding or a grant is secured	Engineer 1 from April 2027; engineer 2 in 2028 (on-premise); then at €500k and €900k recurring revenue
Business developer	Pilots, sales, partnerships	To be recruited with the AI accelerator once funding is secured; not a co-founder	Full-time, once recurring revenue reaches €75k (2028 in the base case)
Head of Commercial (search)	Sales leadership	Has sold to newsrooms, PR agencies or public-integrity bodies; salary + commission + 4–6% vested options	Full-time, once recurring revenue reaches €500k (2030 in the base case)
Fátima López (CVC KTO) · NTT Data mentor	Transfer support · mentoring	Incorporation, contracts, investor readiness	Support

Relationships and ability to raise funds. Competitive research funding (REMISS); broadcaster contacts through university contacts, the CVC transfer office and 4YFN 2026; the AI accelerator and NTT Data Deep Tech BCN programme.

4.2 Shareholding structure

A proposal for the founders to discuss, built on one principle: the people working full-time hold most of the company.

Holder	Proposed at incorporation	Indicative after the €300k pre-seed (€2.25M cap)
Arka Ujjal Dey (full-time)	64%	56.5%
Ernest Valveny (part-time)	10%	8.8%
Dimosthenis Karatzas (part-time)	2%	1.8%
Artemis Llabrés Brustenga (part-time)	2%	1.8%
Employee option pool (incl. Head of Commercial 4–6%, business developer, engineers)	15%	13.2%
UAB / CVC (IP licence)	7% (to negotiate)	6.2%
Pre-seed (€300k, 2027)	—	11.8%

Indicative: one €300k pre-seed as a convertible or participative loan with a €2.0–2.5M cap (≈12% at €2.25M), matched by €485k of ENISA loans and grants. The commercial lead joins as Head of Commercial with salary, commission and 4–6% options from the pool, not a co-founder block. This split is a proposal to confirm in the founders' agreement. Angels are chosen for alignment with what we are building: patient capital (7–10 years, a profitable company is a good outcome), information rights only (quarterly report, annual meeting), no board seat, at most a 1x non-participating liquidation preference and no forced sale without a supermajority. Revenue-based financing is considered only after 2029 and only for marketing. Four-year vesting with a one-year cliff for all founders. Founders keep the majority, as ACCIÓ Startup Capital requires. Have terms checked by a startup lawyer.

4.3 Organisation chart




Figure 6. Current structure (Q4 2026).




Figure 7. Planned organisation, 2031 (base case).

4.4 Positions and profiles

Position	Profile	Hire
Head of Commercial / Business Development	Has sold to newsrooms, PR agencies or public-integrity bodies; leads sales; base + commission + 4–6% vested options	At €500k recurring revenue (2030 in the base case); keep a shortlist warm
Business developer	Runs pilots, demos and follow-up; commission on first-year contracts	At €75k recurring revenue (2028), with the accelerator
ML / software engineers	Retrieval, LLM systems, video, security	1 (≈April 2027) → 2 (2028, on-premise) → 3 at €500k → 4 at €900k (2031); a 5th only once profitable and above €1.5M
Customer success / deployment	Onboarding; offline installs for public-integrity units	At €250k recurring revenue (2029); second at €900k (2031)
Account executives	B2B sales to PR and public sector in the EU and UK	At €900k recurring revenue (2031); a second only once profitable and above €1.5M

4.5 Incorporation plan (headcount)

FTE at year end	2027	2028	2029	2030	2031
Management (Founder & Project Lead, Head of Commercial)	1	1	1	2	2
Business developer	0	1	1	1	1
Engineering and research	1	2	2	3	4
Customer success / deployment	0	0	1	1	2
Account executives	0	0	0	0	1
Total FTE	2	4	5	7	10

Part-time scientific co-founders are not counted as FTE and are not paid in the base case.

4.6 Remuneration system

Role	Gross salary, full year (€k)	Variable / incentive
Founder & Project Lead	38 (2027); 48 until recurring revenue passes €300k (2028–29); then 60–65	Founder equity
Head of Commercial	65 (joins at €500k recurring revenue, 2030)	4–6% options (4-year vesting); variable up to 20% of salary on new ARR
Business developer	39 (joins at €75k recurring revenue, 2028)	Commission ≈8% of first-year contract value; options
Engineers	40 at entry (first hire ≈April 2027) → 46 by 2031 (average per FTE)	Options
Customer success	≈46 (from 2029, at €250k recurring revenue)	Options
Account executives	≈38 base (from 2031, at €900k recurring revenue)	Commission on first-year contract value; on-target earnings ≈€55k

Loaded cost = gross salary × 1.31 ([employer social security ≈31%](#)), as in the Staff expenses tab of the workbook. Customer-success and deployment staff are counted in cost of revenue (SaaS practice). Sales commissions are budgeted at 4% of sales (sales and marketing); payment fees at 2% (cost of revenue). Option pool 15%, 4-year vesting, 1-year cliff. Founder salary: €38k in 2027 and €48k until recurring revenue passes €300k, then €60–65k.

5. SWOT Analysis and Contingency Plans

5.1 PESTEL (external factors)

- **Political** — elections ([Spain 2027](#), EU and UK by 2029) raise demand; any political use must stay neutral.
- **Economic** — newsroom budgets are tight; PR and public-sector budgets are larger but slower or less loyal.
- **Social** — falling trust in information and in AI answers raises the value of visible sources.
- **Technological** — AI-generated images and video increase volume; hosted AI models keep getting cheaper, but web search is ~90% of our cost and GPU rental is rising: H100 one-year rates went from \$1.70 to \$2.35 per GPU-hour between October 2025 and March 2026 ([SemiAnalysis, via IntuitionLabs](#)).
- **Environmental** — offline packs and local models reduce repeated calls and energy use.
- **Legal** — GDPR (political opinions are special-category data), EU AI Act (biometric identification is high-risk or prohibited), copyright and press-content limits on storing excerpts, search-engine terms on storing and reusing results.

5.2 SWOT matrix and strategies

SWOT and contingency plans	Threats	Opportunities
	1. Monitoring suites add similar features. 2. Cheap AI fact-check tools set low price expectations. 3. A wrong label shown in public. 4. Dual use: PR or governments targeting critics. 5. Search API price, quota or policy changes (Google grounded search).	1. Viral images and video need provenance tools. 2. PR teams pay for monitoring already. 3. Elections 2027–2029. 4. EU rules favour auditable tools. 5. Grants for fact-checking and research.
Strengths 1. Working product across text, images, video. 2. Fine-grained labels with quotes and reasons; no verdicts. 3. Published method; source direction 93–96% vs reference on typed claims (provisional). 4. Costs measured per check; usage ≤15% of every plan's price; offline checks €0.01.	Defensive strategies 1. Every label shows who made the call. 2. Terms of use; no person-tracking outside legal mandates. 3. Brave-first search and offline packs to reduce dependence on Google; local judge for sovereignty.	Offensive strategies 1. Pilots with newsrooms for proof and data. 2. PR monitoring as the fastest revenue. 3. Public-integrity units once the chain of custody is ready.
Weaknesses 1. No pilot or paying customer yet. 2. The founder carries sales until about 2030 (Head of Commercial at €500k recurring revenue). 3. English-only interface; no two-factor login. 4. Video claims: wrong side 26–33%; videos ~1.5× their length; local judge 54%. 5. No patent (application rejected); IP licence not yet negotiated.	Survival strategies 1. Business developer at €75k recurring revenue; every hire tied to revenue and a cash floor. 2. €300k from aligned angels plus grants; no dependence on later rounds. 3. Close IP licence and founders' agreement at incorporation.	Reorientation strategies 1. If newsrooms will not pay, lead with PR. 2. If police procurement stalls, sell to electoral and integrity bodies. 3. June 2027 go / no-go review.

5.3 Contingency plans

Risk	Early warning	Plan B	Financial effect
Pilots show low use	Pilot metrics below thresholds (activation <70%, week-2 retention <50%)	Fix workflow issues and rerun; narrow to the segment with the best results	Revenue starts later
Newsrooms will not pay	<3 paid teams by June 2027	Lead with PR; offer newsrooms grant-funded plans	Lower newsroom line
PR agencies convert slower than planned (48 by 2031)	<3 paying agencies by June 2027, or the entry tier needs board approval	Lead with newsrooms and public bodies; delay account-executive hires	Toward the downside case
PR teams need features we lack	Claim tracking and alerts not ready by Q2 2027	Manual weekly reports as a service while building	Slower PR ramp
Wrong label in public	Customer complaint	Labels always show quote and who decided; correction process	Reputation
Police sales need what we lack	Local judge <75% or legal review negative	Sell to electoral and integrity bodies; delay police line	Downside case
Search costs rise or Google limits us	Cost per check +50%, or Google quota not raised	Brave-first search; offline packs; raise credit prices; negotiate a quota	Gross margin 69% → 63% (2031); break-even still 2029
Usage per credit higher than planned	Usage cost above 25% of a plan's price	Enforce credit weights and per-check cost caps; economy search; topic packs; price credit packs	If every credit costs 2× plan: gross margin ~61%, break-even 2030
Pre-seed round slow	No commitments by March 2027	ENISA and ACCÍO first; delay second hire	Lower 2027 burn
NEOTEC rejected	CDTI decision (2028)	Hire the second engineer only once a public body commits to an on-premise pilot; apply for Torres Quevedo	Break-even 2030; lowest cash ~€190k with the second engineer delayed (~€138k if hired on plan)
Incorporation delayed	IP licence not agreed by January 2027	Paid pilots as CVC research contracts; angels sign conditional commitments; prepare grant files in advance	Revenue, ACCÍO grant and hires move back one quarter

6. Operational Plan (Implementation)

6.1 Challenges

- **Prove use** — paid design-partner pilots from November 2026, judged on instrumented metrics.
- **Make it sellable** — Spanish/Catalan interface, check credits and per-check cost caps, claim tracking and alerts (2027), two-factor login; better video accuracy and speed (human-labelled video set, Brave-first search, time budgets).
- **Build the commercial team as revenue allows** — business developer at €75k recurring revenue (2028), Head of Commercial at €500k (2030); until then the founder leads sales with the accelerator's support.
- **Prepare the public-integrity line safely and on time** — legal review, chain of custody and local judge, with the on-premise version ready in 2028 (police and government bodies have shown active interest).
- **Finance** — incorporate around January 2027 (the date depends on the university's spin-off process, which is slow), IP licence, €300k pre-seed round, grants. The next 12 weeks are planned in detail in the final section.

6.2 Schedule




Figure 8. Five-year milestone schedule.

Period	Technical	Commercial	Financial	People
Q4 2026	ES/CA interface; claim fix; check credits and cost caps; new benchmark	Paid design-partner pilots from November; betevé / TV3; PR discovery calls	IP and founders' agreements drafted; aligned-angel list	Accelerator team building: engineer profile; business-developer shortlist
2027	Claim tracking and alerts; 2FA/SSO; event mode; API beta	First paid newsrooms and PR clients; May elections package	Incorporation (target Q1, depends on the university); €300k pre-seed; ENISA €75k; ACCIO	Founder + first engineer; founder leads sales
2028	On-premise version: chain of custody; local judge ≥75%; security audit, ENS	20 newsrooms, 10 PR agencies; first public-integrity licence	NEOTEC; ENISA €100k; revenue €253k	4 FTE (business developer, second engineer)
2029	Video copy detection; echo detection	35 newsrooms, 20 PR, 3 public-integrity units	Revenue €570k; EBITDA break-even (€50k)	5 FTE (customer success)
2030	Offline library growth	55 newsrooms, 32 PR, 6 units	Revenue €946k; EBITDA €6k (hiring)	7 FTE (Head of Commercial, third engineer)
2031	EU languages; API growth	80 newsrooms, 48 PR, 10 units	€1.6M ARR; EBITDA €196k (13%)	10 FTE

6.3 Collaborating entities

Entity	Role	Status
CVC / UAB and CVC transfer office	Research base, IP owners, interim contracting	Active
AI accelerator (4th edition)	Team building (engineer and business developer after funding), mentoring	Active
NTT Data Deep Tech BCN	Mentoring, investor readiness	Active
betevé, TV3/CCMA	Broadcaster pilots and references	Meetings October 2026
Fact-checkers (Verificat, Maldita.es, Newtral, EFE Verifica, Full Fact)	Pilot users, design partners, evaluation	To approach Q4 2026
PR agencies and public-campaign managers (Spain, UK)	Pilot users and first PR customers	To approach Q4 2026
OpenAI, Google (grounded search, Vision), Brave Search	AI judge, web and reverse image search	Commercial; terms of use under legal review

7. Economic-Financial Plan

7.1 General considerations

- Figures in € thousands, constant prices, VAT excluded; salaries rise ~3% a year. Structure follows the NTT Data template; the companion workbook financial_model_v30.xlsx has every formula, a scenario switch and a SaaS-metrics tab.
- Year 0 (Q4 2026): activity at CVC/UAB, on the CVC GPU, by the founders; pilots are paid and contracted through CVC. The company's accounts start in 2027; the first engineer joins after the pre-seed round (~April 2027), the second in 2028 for the on-premise version, and every later hire when its revenue trigger is met and at least three months of costs remain in the bank (\$4.5).
- Revenue on average active customers; most customers prepay annually (8% off), so customers pay in about 20 days on average; suppliers are paid in 30 days. Paid pilots are counted only for customers who convert (pilots that do not convert are upside), with a 25% delivery cost.
- Cost of revenue follows SaaS practice: (1) usage — web search and AI calls at list prices, image APIs and topic-pack refresh — as a share of each line's revenue: newsrooms 15%, PR 15%, public-integrity 10%, API 25%, above the measured 12–14% at full allowance use (\$2.3); (2) production hosting: the CVC GPU in 2027, then cloud GPUs (~€216–501 a month each), database and web, €2k rising to €34k a year; (3) customer-success and deployment staff; (4) a 2% royalty to UAB/CVC (to negotiate) and 2% payment fees. Sales commissions (4%) are a sales and marketing expense; appliance hardware is bought by the client and is not revenue. Gross margin: 77% (2027), 74% (2028), 65%–69% from 2029 (customer success joins in 2029).
- Grants: ACCIÓ Startup Capital ~€90k (2027–2028); NEOTEC €220k (2028–2029, cash 2028 and 2030). ENISA loans €75k (2027) and €100k (2028).
- [Corporate tax 15% in the first two profitable years, 25% after](#), once prior losses are used. No dividends.

7.2 Investment estimation

	2027	2028	2029	2030	2031
Reference appliance for public-integrity pilots (L40S server, ~€15–18k; estimate, no vendor quote yet)	0	17	0	17	0
Laptops and equipment	8	3	5	5	6
Office fit-out	0	0	6	0	0
Total investment	8	20	11	22	6
Amortization (4 years)	2	7	10	15	15

7.3 Revenue estimate

Customer numbers and revenue by line are in §2.4. The most important assumptions, and how the pilots will test them:

Assumption	Base case	Evidence today	Test
Newsroom teams pay	8 (2027) → 80 (2031) at €5.7k (€5.2k after the prepay discount)	None yet; self-serve tools are cheap	Pilot conversion and willingness to pay, Q4 2026 – Q1 2027
PR clients	3 (2027) → 48 (2031) at ~€15.5k (€1,200/month + crisis add-ons; €14.3k after the prepay discount)	None yet; two-month paid pilots test it	5 discovery calls; pilot budget question; crisis pilot
Public-integrity licences	First in 2028 → 10 (2031) at ~€30k (cloud €15–25k, on-premise €35–50k)	Needs chain of custody and local judge	Legal review; 2 exploratory talks
Costs per check	€0.09 typed claim, €0.17 image post, €0.07–0.11 per video claim checked (list prices); €0.01 offline	Spend ledger, 12–18 Oct 2026 (one team, development use)	Spend ledger per customer in the pilots; Brave-only coverage test
Usage within allowances	600 credits a month per newsroom plan, 1,000 per PR agency; usage ≤15% of price at full use	None yet	Pilot usage logs; credit metering (Q4 2026)

7.4 Staff and other expenses

Staff incl. social security (€k)	2027	2028	2029	2030	2031
Founder & Project Lead (founder)	50	63	63	79	85
Head of Commercial (at €500k recurring revenue)	0	0	0	85	86
Business developer (at €75k recurring revenue)	0	51	52	54	54
ML / software engineers	39	121	123	189	257
Customer success / deployment (in cost of revenue)	0	0	62	63	128
Account executives	0	0	0	0	52
Total staff	89	235	300	470	662

Other Expenses (€k)	2027	2028	2029	2030	2031
Rent	4	8	16	24	30
Software and internal IT (production hosting is in cost of revenue)	6	7	8	10	12
Lawyers (incorporation, IP licence, contracts, GDPR / AI Act, search-engine terms)	22	12	12	14	14
Mobile phones	1	2	3	4	5
Accounting & gestoria	5	6	8	9	10
Marketing & travel (8% of revenue, €15–90k)	15	20	46	76	90
Insurance & others	3	5	7	9	10
External support (security audit, ENS / ISO 27001, legal review)	5	25	30	20	20
Subscriptions (data, associations)	3	5	7	9	10
Sales commissions (4% of sales)	4	10	23	38	60
Total other expenses	68	100	160	213	261
R&D (annotation, evaluation, audits)	10	25	35	40	45
Memo: production hosting, cloud GPU, database (in cost of revenue)	2	15	22	28	34

7.5 Profit and loss

P&L (€k)	2027	2028	2029	2030	2031
Newsrooms	60	99	196	262	393
PR & communications	21	93	215	372	572
Public-integrity / OSINT	0	20	70	150	260
API	0	15	50	110	200
Paid pilots	12	26	39	52	72
TOTAL SALES	93	253	570	946	1,497
Usage (web search, AI judge, topic packs) and pilot delivery	(15)	(41)	(90)	(151)	(239)
Production hosting, cloud GPU, database	(2)	(15)	(22)	(28)	(34)
Customer success and deployment staff	0	0	(62)	(63)	(128)
Royalty UAB/CVC (2%)	(2)	(5)	(11)	(19)	(30)
Payment fees (2%)	(2)	(5)	(11)	(19)	(30)
GROSS MARGIN	72	187	374	666	1,036
Gross margin %	77%	74%	66%	70%	69%
Grant income	45	155	110	0	0
R&D expenses	(10)	(25)	(35)	(40)	(45)
Staff expenses (excl. customer success)	(89)	(235)	(238)	(407)	(534)
Other expenses (incl. 4% sales commissions)	(68)	(100)	(160)	(213)	(261)
EBITDA	(50)	(18)	51	6	196
EBITDA before grants (operating)	(95)	(173)	(59)	6	196
EBITDA margin	n/m	-7%	9%	1%	13%
Amortization	(2)	(7)	(10)	(15)	(15)
EBIT	(52)	(25)	41	(9)	181
Interest (ENISA)	(2)	(8)	(11)	(10)	(8)
EBT	(54)	(33)	30	(19)	173

2027	2028	2029	2030	2031	
Income tax	0	0	0	(15)	
NET PROFIT	(54)	(33)	30	(19)	158

Gross margin 77% of sales in 2027, 74% in 2028 and 65%–70% from 2029, after hosting, customer success and pilot delivery. EBITDA turns positive in 2029 (€51k, including €110k of NEOTEC grant income; before grants it is –€59k, so operating break-even is 2030), is about flat in 2030 (€6k) while the Head of Commercial and a third engineer join, and reaches €196k (13% of sales) in 2031. Losses of €87k in 2027–2028 are offset against later profits; income tax starts in 2031.

7.6 Balance sheet

Balance sheet 2021	2027	2028	2029	2030	2031
Net fixed assets	6	19	20	27	18
Clients	4	12	26	43	67
Grants receivable	0	0	110	0	0
Treasury	329	383	304	358	495
TOTAL ASSETS	339	414	460	428	580
Social capital (incl. share premium)	310	310	310	310	310
Reserves	0	(54)	(87)	(57)	(76)
Net profit	(54)	(33)	30	(19)	158
EQUITY	256	223	253	234	392
ENISA loans	75	175	175	150	125
Trade creditors	8	16	32	44	63
TOTAL LIABILITIES AND EQUITY	339	414	460	428	580

7.7 Cash flow

2027	2028	2029	2030	2031	
EBITDA	(50)	(18)	51	6	196
Financial result	(2)	(8)	(11)	(10)	(8)
Income tax	0	0	0	0	(15)
CURRENT CASH FLOW	(52)	(26)	40	(4)	173
Change in working capital	4	0	(108)	105	(5)
GROSS CASH FLOW	(48)	(26)	(68)	101	168
Investment	(8)	(20)	(11)	(22)	(6)
FREE CASH FLOW	(56)	(46)	(79)	79	162
Capital increases	310	0	0	0	0
Loans drawn / (repaid)	75	100	0	(25)	(25)
FINANCIAL CASH FLOW	385	100	0	(25)	(25)
TREASURY VARIATION	329	54	(79)	54	137
ACCUMULATED CASH	329	383	304	358	495




Figure 9. Lowest year-end cash is €304k (2029), about 5.8 months of that year's costs; with hiring tied to revenue, the downside case keeps at least €166k.

7.8 Profitability: VAN (NPV) and TIR (IRR)

€k	2027	2028	2029	2030	2031	Terminal value
Free cash flow	(56)	(46)	(79)	79	162	5,988 (4× 2031 sales)
Discount factor (30%)	0.769	0.592	0.455	0.350	0.269	
Present value	(43)	(27)	(36)	28	44	1,613

VAN (NPV) at 30% = €1,578k; TIR (IRR) = 199%, including a terminal value of 4× 2031 sales. Without the terminal value, the five-year cash flows give NPV €(35)k. These are the project's returns, driven by the terminal value; they are not an investor's multiple.

NPV (€k) assumptions	Exit 3× sales	4×	6×
Discount 25%	1,442	1,933	2,914
Discount 30%	1,175	1,578	2,384
Discount 35%	963	1,296	1,964

7.9 Financing plan

Source	€k	Timing	Conditions
Founders	10	Q1 2027	Equity
Pre-seed from aligned angels (media, publishing, fact-checking, PR operators, impact investors; BStartup, ESADE BAN, IESE BAN)	300	Q1–Q2 2027	Convertible loan, €2.0–2.5M cap (indicative); information rights only, no board seat
ENISA young-founder loan	75	2027	Own funds ≥50% of the loan, paid in first
ACCIÓ Startup Capital	≈90	2027–2028	Up to €99k, 75% of costs ; company ≤18 months; founders hold majority
Optional later round (not planned)	—	From strength, e.g. >€600k ARR	Only to grow faster; aligned investors; founders keep control
ENISA loan	100	2028	Own funds ≥100% of the loan
CDTI NEOTEC	220	2028–2030	Non-dilutive (CDTI; up to €250k, 70% of the budget)
Institutional seed	—	—	Not planned (Option 1: high-control, capital-efficient)

	Amount	%
Sources		
Equity (founders €10k, pre-seed €300k)	310	39.0%
ENISA loans	175	22.0%
Grants (ACCIÓ Startup Capital, NEOTEC)	310	39.0%
Total sources	795	100.0%
Uses		
Operating losses before grants, 2027–2029	348	43.8%
Working capital at end 2029	104	13.1%
Capital expenditure 2027–2029	39	4.9%
Loan repayments	0	0.0%
Cash at end 2029 (buffer)	304	38.2%
Total uses	795	100.0%

7.10 Scenarios

Scenario	Assumptions	2031 sales	2031 EBITDA	EBITDA ≥ 0	Lowest year-end cash*
Downside	About half the customers, a year later (45 newsrooms, 24 PR, 5 units); hiring waits for revenue	773	72	2031	166
Base	80 newsrooms, 48 PR, 10 units, API €200k; list search prices	1,497	196	2029	304
Base, NEOTEC rejected	Same customers and team; the €220k NEOTEC grant is not awarded	1,497	196	2030	138
Upside	115 newsrooms, 75 PR, 14 units, API €300k	2,251	769	2028	351

*Lowest year-end cash, 2027–2031 (€k). In the downside case, hiring triggers are met later, so the team stays at 3–4 people and cash never falls below €166k (computed with the plan model). If, in addition, NEOTEC is rejected, the second engineer is hired only once a public body commits to an on-premise pilot; cash then stays positive but thin (≈€35k in 2030), which the June 2027 go / no-go review is designed to catch. 'Base, NEOTEC rejected' is the base case without the grant (workbook): cash falls to €138k in 2030 and recovers.

Unit-cost sensitivity (base-case customers and prices; only the cost of each check changes):

	Assumption	2031 gross margin	2031 EBITDA	EBITDA ≥ 0	Lowest cash*
Economy search	Brave-first search holds coverage: usage shares 5% newsrooms, 6% PR, 10% API	77%	317	2028	335
Search prices +50%	Google or Brave raise prices, or quotas tighten: usage shares ×1.45	63%	107	2029	249
Cost per credit ×2	Every credit costs twice the planning figure (harder claims, more search rounds)	61%	129	2030	186

*As above. Gross margin stays above 60% in every case with hosted AI models. Search pricing and credit discipline matter far more than the choice of AI judge, which is ≈8% of the cost of a check; an on-premise judge is upside, not plan.

7.11 Investor metrics and use of funds

The SaaS metrics investors compare (SaaS metrics tab of the workbook). Churn — 15% a year for newsroom plans, 20% for PR clients, 10% for public-integrity licences — and customer-acquisition cost are assumptions until the first contracts; customer acquisition cost includes all sales and marketing, including the Head of Commercial and the founder's selling time (50% of the founder's cost until the Head of Commercial joins, 25% after). The metrics are illustrative until the first contracts; the Rule of 40 is shown from 2029, when the revenue base is meaningful.

2026	2027	2028	2029	2030	2031
ARR at year-end (€k)	85	293	609	1,036	1,606
ARR growth	—	245%	108%	70%	55%
Gross margin	77%	74%	66%	70%	69%
EBITDA margin	n/m	-7%	9%	1%	13%
Rule of 40 (revenue growth + EBITDA margin, from 2029)	—	—	134%	67%	71%
Customers at year-end (excl. API)	11	31	58	93	138
Blended customer acquisition cost (€k)	4.0	5.2	4.8	6.1	6.0
CAC payback (months, gross-margin basis)	8	9	9	10	10
LTV / CAC	9.9	7.8	8.1	6.9	7.1
Burn multiple (net burn / net new ARR)	0.66	0.22	0.25	—	—
Year-end cash (€k)	329	383	304	358	495

Use of funds and the milestones. The €300k round (with ENISA €75k and ACCIÓ ≈€90k in 2027, then ENISA €100k and NEOTEC €220k) funds the founder and the first engineer from incorporation, the second engineer and the on-premise version in 2028, the Spanish/Catalan interface, credit metering and two-factor login, and the pilots. Later hires are paid from revenue. It must deliver the June 2027 go / no-go proofs: ≥5 paying customers or €60k contracted ARR, ≥50% pilot retention and ≥80% agreement with a human-expert reference. The company reaches EBITDA break-even in 2029 with €609k ARR (2030 before grant income) and needs no further equity. A second round is optional and would be raised only from strength (for example above €600k ARR) to grow faster.

Why the plan is investable. Capital efficiency: €310k of equity plus €485k of non-dilutive finance reaches break-even in 2029 (2030 before grants) with no further round; founders keep control, and every hire is paid for by revenue. Software margins with hosted AI: gross margin 65–77% because usage is sold in credits and measured per check. A defensible niche: no competitor shows, per claim, who confirms and who disputes it with the quote and reason behind each label across text, images and video. Honest risks: no customer yet, video-claim accuracy still below target, and churn and acquisition cost to be proven in 2027.

Annex I – Abbreviated CVs of key people

Arka Ujjal Dey — Founder & Project Lead. Postdoctoral researcher, Computer Vision Center; first author of the RAV study and CRAVE method. Full CV on request.

Ernest Valveny Llobet — Co-founder & Scientific Lead. Professor, UAB; CVC researcher. Full CV on request.

Dimosthenis Karatzas — Co-founder (advisory for now). Professor, UAB; CVC researcher. Full CV on request.

Artemis Llabrés Brustenga — Co-founder (advisory for now). CVC technical staff. Full CV on request.

First engineer after funding (≈April 2027), second engineer in 2028 (on-premise), business developer at €75k recurring revenue — recruited with the AI accelerator; profiles in §4.4.

Annex II – Business Model Canvas

Block	Content
Key partners	CVC/UAB; AI accelerator; NTT Data; fact-checkers; broadcasters; OpenAI, Google (search, Vision), Brave Search.
Key activities	Engine R&D; labelling quality; offline packs; pilots; customer onboarding.
Key resources	Product v7; label scheme and expert-labelled data; offline evidence library; research team; licensed software and know-how.
Value propositions	Who confirms a claim, who disputes it and what is left out — sourced, dated, labelled, no verdict — for text, images and video — replacing hours of manual investigation.
Customer relationships	Paid pilots; team plans; agency workspaces with per-crisis add-ons; cloud and on-premise licences with installation and training.
Channels	Direct sales; pilots; fact-checking networks; PR agency contacts; events (4YFN); accelerator.
Customer segments	Newsrooms and fact-checkers; PR and communications teams; public-integrity and police OSINT units; AI product teams.
Cost structure	Salaries; web search (≈90% of the variable cost) and AI calls; cloud GPU and hosting; security and legal; royalty or equity to UAB/CVC.
Revenue streams	Paid pilots; monthly and annual plans (8% off for prepayment); per-crisis add-ons; event packages; custom topic packs; cloud and on-premise licences and set-up fees; API usage.

Annex X – Other relevant information

- Product one-pager, client deck and quick-start guide (counterview v7, October 2026).
- RAV report and CRAVE preprint ([arXiv:2504.10166](https://arxiv.org/abs/2504.10166)).
- Market study for counterview by Viromii for CVC, 19 March 2026 (file “2ndMarket Study_Counterview_CVC_19032026.pdf”); market sizes on p. 26.
- UAB/CVC invention disclosure “CVRAV”, 7 March 2025 (available on request; contains personal data).
- Unit economics, accuracy and performance reference (cost_acc_answers, 13 October 2026; spend ledger and accuracy runs, 12–18 October 2026; internal).
- External pricing review, round 1, and our response (feedback_round1.txt, 20 October 2026).
- Contacts: auzey@cvc.uab.cat · ernest@cvc.uab.cat

Immediate action plan: the next 12 weeks

From **12 October 2026 to 3 January 2027**. Purpose: get real users on the product, measure whether they find it useful and would pay, and be ready to incorporate in January. Owners are proposals, confirmed at the first Friday review. A 30-minute review every Friday tracks the targets below, using the *Pilot KPIs* tab of the financial workbook.

Targets for week 12

Area	Target	Why it matters
Pilots	8 paid pilots started (4 newsrooms/fact-checkers, 3 PR, 1 public-integrity); ≥6 completed	Real use, not opinions
Pilot metrics	≥70% activated, ≥60% useful checks, ≥50% week-2 retention, UMUX-Lite / SUS ≥68	Already instrumented
Willingness to pay	Answers from every pilot team; ≥2 pilots convert to paid plans or sign letters of intent	Tests the price hypotheses
Product	Spanish/Catalan interface; unresolved claims 16% → <8%; PR alerts prototype; default search path chosen (Brave-first if coverage holds); check credits and per-check cost caps live	Must-haves before selling
Quality	New benchmark results (AVeriTeC, WiCE, 200-claim news set) with a human-expert reference; source direction ≥90%, finer labels ≥80%	Replaces Perspectrum
People	Engineer and business-developer profiles agreed with the accelerator and a shortlist ready, to hire once funding closes; business-developer shortlist ready for €75k recurring revenue	Commercial gap
Company	IP licence terms and founders' agreement drafted; incorporation papers ready; list of 20 angels	Incorporate around January

Phase 1 — Launch pilots (weeks 1–3, 12 Oct – 1 Nov)

Week	Action	Owner	Done when
1	With the accelerator: define the engineer profile (hire after funding) and the business-developer profile (hire at €75k recurring revenue)	Arka, Ernest	Profiles and shortlist ready
1	Confirm betevé and TV3/CCMA meeting dates; show the recorded-video check; invite them into the pilot	Arka, Fátima	Dates set; invitation sent
1–3	Recruit the pilot cohort: newsrooms and fact-checkers (Verificat, Maldita.es, Newtral, EFE Verifica, Full Fact), Spanish and UK PR agencies, one public-integrity body	Arka, Ernest	8 teams confirmed
1–3	Spanish and Catalan interface	Arka	Shipped before broadcaster use
2	Add willingness-to-pay questions to the pilot feedback rounds (price ranges per package)	Arka	Questions live
2–3	Choose the replacement benchmarks; start labelling the 200-claim news test set	Ernest (Dimos, Artemis advising)	Set defined; labelling started
3	Pilots start (paid: newsrooms 1 month €500, PR 2 months €2,500), contracted through CVC	All	First teams active

Phase 2 — Learn from use (weeks 4–6, 2 – 22 Nov)

Week	Action	Owner	Done when
4–6	Weekly pilot check-ins; track activation, useful-check rate, retention	Arka	KPIs updated every Friday
4–6	Claim decontextualisation (resolve "he", "this deal" before search)	Arka	Unresolved claims below 8%
4–5	Fact-check badge (outlet's own rating shown with the source)	Arka	Shipped
4–5	Cost and speed tests: Brave-only vs Brave + Google coverage on 20 claims (≈€2); a 20-minute video on today's setup (≈€15); batched local-judge test (no API cost); one week of monitoring to measure evidence reuse	Arka, Ernest	Default search path and video allowance decided
4–6	5 discovery calls with PR agencies and public-campaign managers (Spain, UK): hours spent per claim today, crisis-investigation needs and prices	Arka, Ernest	Notes on needs and prices
5	IP meeting with CVC transfer office and UAB: licence of software and know-how (no patent); equity rather than royalty	Ernest, Fátima	Proposed terms on paper
6	Founders: roles, time commitment, cap-table proposal (\$4.2)	All founders	Agreement in principle

Phase 3 — Convert and design (weeks 7–9, 23 Nov – 13 Dec)

Week	Action	Owner	Done when
7	Pilot results: metrics, feedback, willingness to pay per segment	Arka, Ernest	Results written
7–9	Convert the best pilot teams to paid plans from January	Arka	≥2 accepted or letters of intent
7–9	PR claim-tracking and alerts prototype (daily re-run, alert on new disputing outlets)	Arka	Prototype shown to PR pilots
8–9	Public-integrity line: 2 exploratory talks (EU/UK); scope a legal review (AI Act, GDPR) — no selling	Ernest, Arka	Requirements listed
8–9	Business developer: shortlist 2–3 candidates for 2028; keep potential Head of Commercial candidates warm	Arka, mentor	Shortlist of 1–2

Week	Action	Owner	Done when
9	Check ENISA, ACCIÓ Startup Capital and NEOTEC requirements	Fátima, Arka	Checklist ready

Phase 4 — Decide and prepare 2027 (weeks 10–12, 14 Dec – 3 Jan)

Week	Action	Owner	Done when
10	Review all week-12 targets; decide the lead segment for Q1 2027	All founders	Decision recorded
10–11	Update the business plan (v31) and workbook with pilot data, prices heard and real costs	Arka	v31 ready
10–11	Benchmark results on the new sets; first paper drafts (labels and "who said it")	Research team	Results and drafts
11–12	Incorporation documents, founders' agreement and IP terms ready for signature	Fátima, Ernest	Ready for signature (date depends on the university)
11–12	Aligned-angel list (media, publishing, fact-checking and PR operators, impact investors; accelerator, NTT Data, CVC/UAB, BANC, BStartup, ESADE BAN, IESE BAN) and founder-friendly convertible terms with a lawyer	Arka, mentor	List and terms ready
12	4YFN 2027 application	Arka	Submitted

June 2027 go / no-go

counterview continues as a company after June 2027 only if: (1) the commercial plan is staffed: a business developer is shortlisted to join once recurring revenue reaches €75k; (2) at least 5 paying customers, or €60k of signed annual contract value; (3) pilot users keep coming back ($\geq 50\%$ still active at the end of their pilot); (4) source labels agree with a human-expert reference at least 80% of the time on the new benchmark set, video claims included; (5) measured usage cost stays within the model ($\leq 15\%$ of newsroom revenue). If not, we license the technology or stop.