FYYAO

PRE-SEED ROUND · OCTOBER 2026

The agent that pays for and manages your digital life

> **"I need a video tool for my business, USD 30 a month at most."**
>
> "I compared 9 options. I recommend this one: USD 24 a month, it exports without a watermark. The two cheapest ones do not have the feature you asked for. Shall I buy it?"

One account · one agent · zero invisible spend

US + Brazil · Raising USD 700,000 for 14 months

Founders: Dany Meza · Alejandro Reyes

THE PROBLEM

People run their digital economy by hand — and it costs them

Nobody teaches you to manage this. And the questions have nowhere to be answered:

- What should I buy for what I need?
- How much should I be paying?
- Is there something better for less?
- **Am I still using this thing I pay for every month?**
- When does it renew?
- Am I paying for two services that do the same thing?

The result, measured in money:

- **8 to 20 subscriptions** active per person or business, with no inventory and no owner.
- Between **8% and 12% of digital spend is waste**: duplicates, unused seats, badly chosen plans.
- In **Brazil**, paying for international services adds FX spread, IOF tax and cards the provider rejects.
- And nobody cancels because of the hassle, not for lack of wanting to.

The 8–12% range has to be validated with our own data: it underpins 30% of revenue.

WHY NOW

Three things made this possible in the last 18 months

① Agents learned to use tools

A model can now look things up, compare and act reliably. Two years ago this was a demo; today it is infrastructure.

② AI spend became invisible

Every person and every team bought their own tools. Nobody knows how much is being spent. "Am I still using this?" went from curiosity to budget line.

③ The international rail stopped being expensive

Settling through new rails and going local lets us charge **5% in Brazil where the bank charges 8–15%**.

> **The window:** whoever first combines **agent + permissions + rail + history** keeps the user's financial relationship. That position is not recovered afterwards.

THE SOLUTION

One account. One agent with permissions. Your digital spend under control.

1. **The agent finds and decides.** You tell it what you need and how much you want to spend. It searches, compares, recommends and explains why.
2. **The agent actually pays.** It is not an adviser that hands you a link: it executes the payment with the funds in your account, with the rate and the cost in plain sight.
3. **The agent manages and trims.** It spots renewals, price rises, unused services and duplicates. It proposes and executes cancellations.
4. **You keep control.** Permissions by category and by amount, confirmation above the threshold you set, and a full record of every action.
5. **The account holds it all together.** Balance, international payments, currency exchange and — in Brazil — paying without depending on an international card.

HOW IT WORKS

The flow

NEED → **AGENT** → **SEARCH** → **COMPARISON** → **RECOMMENDATION** → **AUTHORISATION** → **PAYMENT** → **RECORD** → **MANAGEMENT**

You: "I need an AI tool to make videos for my business, USD 30 a month at most."

Agent: "I compared 9 options. I recommend this one: USD 24 a month, it includes what you need and lets you export without a watermark. The two cheapest ones do not have the feature you asked for. Shall I buy it?"

And every month, just:

"I found an alternative that is 35% cheaper with equivalent features. Want me to switch?"
"I spotted a subscription unused for 62 days. Shall I cancel it?"

Product metrics: recommendations accepted · operations executed without intervention · **savings generated per user**.

PRODUCT

The core is built and verified. What is missing is exactly what this round buys.

What already works

- The agent compares the catalogue, picks **the cheapest one that does the job** and explains why it ruled out the rest.
- **Permission engine:** per-transaction cap, monthly budget, human confirmation, revocation and an audit of **every** decision, including the rejected ones.
- **Double-entry ledger:** persistent, append-only, and retrying a payment does not charge it twice.
- **153 automated tests.** Verified with the real process: a charge is made, the process is restarted, and the balance, the permission and the budget are still there.

What does not exist yet

- **Users and revenue: zero.**
- The real payment rails (today they are test adapters).
- Identity verification and the written legal opinion.
- Signed payment partners.

What this round buys

The capital does not go into buying users: it goes into connecting the real rails, meeting regulation and getting the **first paying users**.

Use of funds	USD	%
Engineering: agent, permissions, ledger	USD 208,000	30.8%
Operations and core team	USD 192,000	28.4%
Compliance, legal and regulatory	USD 134,000	19.8%
G&A	USD 57,353	8.5%
Brazil acquisition	USD 54,000	8.0%
US acquisition	USD 31,000	4.6%
Total 14-month plan	**USD 676,353**	**100%**

MARKET

We start with people who already spend, not people who might

The first user: independent professionals and small businesses that **already pay between USD 100 and USD 250 a month** for software, AI and international services. They have the pain, the budget and the decision in their own hands.

Bottom-up sizing

```
US      professionals and service SMBs with digital spend
        × USD 200/month of managed spend
Brazil  companies in services and technology that pay internationally
        × USD 120/month of managed spend
──────────────────────────────────────────────────────
= target GMV -> revenue = GMV x (rail + savings + subscription)
```

The advantage: it can be counted by hand

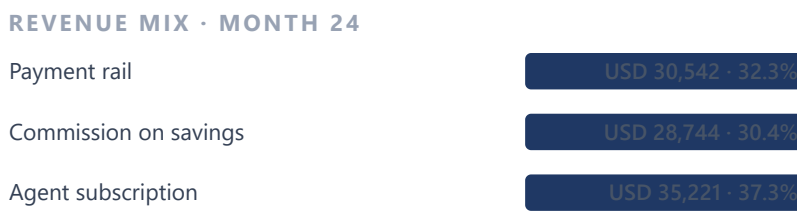
- The launch is not sized by a consultancy: it is built from **a named list of 1,000 target customers**.
- Sources: freelancer and agency communities, accounting firms, coworking spaces and technology associations in both countries.
- **What the model projects is conservative against that:** 5,842 users at month 24 out of a universe of millions.

> The constraint is not the market: **it is CAC**.

BUSINESS MODEL

Three revenue streams per user — and the biggest one isn't the subscription

Revenue	Mechanism	Brazil	US
1 · Payment rail	% of digital spend that runs through the account	**5.0%**	1.5%
2 · Commission on savings	35% of the savings found and executed (10% of spend)	10% × 35%	10% × 35%
3 · Agent subscription	Premium tier per user	USD 5	USD 9
ARPU		**USD 15.20**	**USD 19.00**



The one thing I want you to take away: the agent subscription is not what the business depends on. **If tomorrow a free model did the recommending, we would still be paid for moving money and for saving it.** We are not an AI wrapper: the AI is the interface.

Unit economics

	Brazil	US
Digital spend managed	USD 120/month	USD 200/month
ARPU	**USD 15.20**	**USD 19.00**
Direct cost (rails + compliance + support + fraud)	USD 4.18	USD 4.90
Contribution per user	**USD 11.02 · 72.5%**	**USD 14.10 · 74.2%**
Target CAC	USD 60	USD 110
CAC payback	**5.4 months**	**7.8 months**
Monthly churn	4.5%	3.5%
LTV / CAC	**4.1x**	**3.7x**

The three assumptions to validate

- **CAC (USD 60 / 110):** in consumer, growth is bought. It decides how many users each dollar brings.
- **Digital spend per user (USD 120 / 200):** a third of revenue comes from here.
- **Churn (4.5% / 3.5%):** at 6% the model stops closing.

What I do not promise

- I do not promise a 10x LTV/CAC. These are healthy numbers: payback under 8 months and LTV/CAC between 3.7 and 4.1.
- The metric I lean on today is **payback**, not LTV.
- The job of this round is to turn assumed churn into **measured churn**.

COMPETITION

Everyone has one piece. Nobody has all four.

Player	What it does well	What it lacks
ChatGPT / Gemini / Claude	They find and recommend, for free	**They do not pay, do not cancel, do not manage. They leave you the link and the work**
Rocket Money / Trim	They proved that people **pay** to have their subscriptions managed (commission on savings)	They are not agents, they have no rail of their own, they do not operate in Brazil
Wise / Payoneer / Airwallex	Cheap, reliable rails	They are plumbing: they do not discover, recommend or manage
Banks and cards	Trust, balance sheet, licences	They do not compete for the decision, they compete for the transaction
Clara / Jeeves / Ramp	Excellent spend management for **companies**	They do not serve the individual or the small business; they do not work with the agent
The real status quo	Zero friction today	**You keep paying for what you do not use, month after month**

> You need **agent + permissions + rail + history** at the same time, and every piece is held by someone different. **Rocket Money is the best proof that this market pays.** Nobody has done it yet with an agent and its own rail in two countries.
>
> **"We do not compete with ChatGPT on the recommendation. We compete on execution."**

STARTING POINT

The product exists. The traction does not. This is the starting point.

Product

Core built and verified.

The agent does not execute real payments yet: the rails are test rails.

Customers

None paying.

No pilots running.

Revenue

Zero.

Not a single real transaction processed.

What exists today

- The **product core working**: agent, permissions and ledger, with 153 tests and verified by restarting the process.
- The thesis and the **full financial model**, with every assumption explicit and checkable.
- The **funding plan**: what we build, with how much and with which decision milestones.

The first three pieces of evidence this round buys

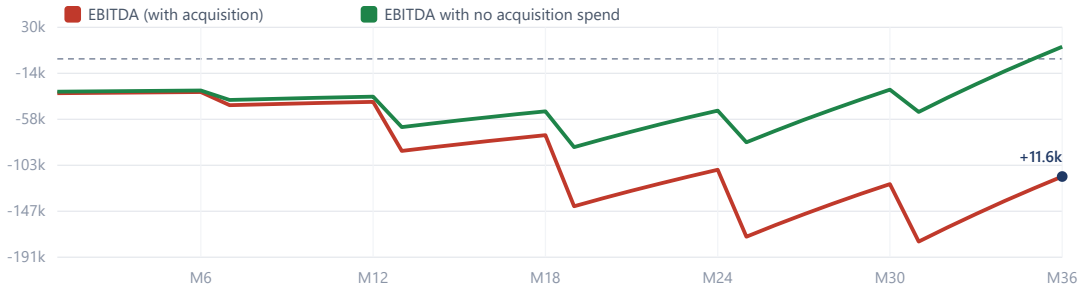
- A product that **executes real payments** in the US and Brazil.
- **Two payment partners** signed, one per country.
- The **written legal opinion** on payments initiated by the agent.

> **Being clear about this is part of the argument:** one paying user is worth more than fifty enthusiastic interviews. The goal of this round is to have the first one — with its CAC and its churn measured.

THE CENTRAL ARGUMENT

Phase one buys product. Phase two buys users. They don't mix.

	Phase 1 · Build	Phase 2 · Growth
What it buys	Product and compliance	Users
Who funds it	**Round 1: USD 700,000**	**Round 2: USD 1,500,000**
Users at the end	1,053	8,706
Monthly revenue at the end	USD 16,976	USD 141,059
Acquisition spend	USD 1,500/month	up to USD 91,000/month
What is measured	That the product pays	**Real CAC and churn**
Deficit of the phase	**USD 631,359**	**USD 2,200,559**



In phase 1 we do not buy a single user too many: acquisition spend is USD 1,500 a month, because there is still no way to charge a real user: the payment rails and identity verification are still in test mode. **And the honest number in the model:** with no acquisition spend, the operation covers its structure only once it reaches ~17,868 users (month 36, +USD 11,628). Before that, every new user is an investment that has to be decided, not a certainty.

Why this separation is the most important thing in the deck: if the two phases are mixed, we are asking for a million and a half to build once, and the investor is right to be wary. By separating them, **each round buys something different and verifiable**: the second is raised on measured data, not assumptions.

Construction milestones, not growth milestones

Milestone	Timing
1 · Partners and legal structure. 2 signed payment-partner quotes (one per country) · written legal opinion on agent-initiated payments	**Months 3-6**
2 · The MVP executes real payments. ~119 pilot users · revenue USD 1,908/month · deficit USD 195,583	**Month 6**
3 · Unit economics measured. ~459 users paying · revenue USD 7,431/month · **CAC and churn measured, not assumed** · deficit USD 454,684	**Month 12**
4 · Round 2. ~1,053 users · revenue USD 16,976/month (ARR USD 0.20M) · this round leaves cash of **USD 68,641**	**Month 14**
5 · Series A. ~8,706 users · revenue USD 141,059/month (ARR USD 1.69M)	**Month 27**
Projection. ~17,868 users · revenue USD 290,341/month (ARR USD 3.48M)	**Month 36**



The two decision markers:
Month 6: if the MVP does not execute real payments in both countries, the product is the problem. We stop and fix it before spending the rest.
Month 12: if measured CAC or churn is over 50% above the assumption, **we do not switch on acquisition** in month 13.

*Why Round 2 is in month 14: that is where the accumulated deficit reaches **USD 631,359** and Round 1 cash is left at **USD 68,641**. **The model says when to raise, not ambition.** New users are capped by the acquisition budget (budget ÷ CAC): we do not project users without the money to bring them.*

Focus is a decision — and we wrote it down

Out of scope for this stage

- ❌ General marketplace
- ❌ Vehicles
- ❌ Credit
- ❌ Insurance
- ❌ Investments
- ❌ Full enterprise system
- ❌ Cryptocurrencies as a product
- ❌ Social network
- ❌ Third-party agent marketplace

Business decisions

- ❌ **We do not resell accounts** for ChatGPT, Claude, Gemini or any provider. The user signs up directly where that applies: we are neutral. That keeps us out of the terms of service, out of account ownership and out of depending on a third party.
- ❌ **We do not take an affiliate commission for recommending.** If the provider paid us, our incentive would stop being to recommend what is best for you. **We are paid through the rail and the savings: that is why the recommendation is credible.**

> A fund prefers a small, executable scope over an ecosystem that cannot be built. And these last two lines are what hold the business asset together: the user's trust.

TEAM

Two founders splitting the problem

Dany Meza · Founder

Vision, product, business model and commercial development. He comes from retail and distribution, not from theory. He ran a minimarket with a butcher's counter and staff reporting to him and, using his point-of-sale data, he found that low-demand hours did not cover their cost: **he moved the opening hours from 07:00–21:00 to 09:00–19:00.** He built and ran Comercial León with several business units — minimarket, butcher's counter and prepared food — and handled wholesale buying, in-house logistics and volume negotiation.

Alejandro Reyes · Technical co-founder

Full-stack development. Platform architecture, **permission engine**, multi-currency ledger and integration of the payment partners.

Why this combination matters

The business we sell is **money and real operations**. A founder who managed margins, inventory, staff and cash flow understands the customer better than one who only knows software. And having development inside the team means that **the product does not depend on hiring a third party**.

A good idea on its own does not build a company: the difference is in executing, selling, controlling costs and understanding the customer.

What this round adds

- **Fintech and regulatory specialist:** payments, compliance, risk and relationships with financial players. It is the most critical gap, because it decides whether we can operate.
- **Development team** to speed up building the platform.
- **Growth** (months 6-12), once there is a product to sell and CAC can be measured.

COMPLIANCE

How we operate legally — and what we won't do

The specific risk of this product

An agent that pays with the user's money is the most delicate combination there is: **a payment initiated by a third party with the user's authorisation**. If the user authorises it and something goes wrong, the bank does not refund. That is why the **anti-fraud engine, the limits and the audit trail are part of the product from day 1**, not a later phase.

How we operate

- **Phase 1 — licensed partners:** custody, rails and settlement. We are not the custodian.
- **Phase 2 — our own licences:** FinCEN/MSB and *money transmitter* at state level in the US; *instituição de pagamento* from the BCB and an eFX structure in Brazil.
- **Real budget: USD 360,000 over 24 months** (advisory, licences and bonds). It is in the model.
- **KYC/KYB and continuous monitoring:** USD 1.20 per user in Brazil and 1.80 in the US, already in the unit economics.

What we are NOT going to do

- ❌ We will not be a bank or a custodian with our own balance sheet.
- ❌ We will not lend or sell crypto as a product.
- ❌ **We will not present "the agent pays on its own" as a launch feature.** It launches with explicit authorisation and thresholds; automation is earned with history.
- ❌ We will not promise to operate in more countries than we can comply in.

The five risks that can kill this

Risk	How we test it cheaply	If it fails
CAC above USD 60/110	Test of 4 channels on a limited budget, measuring CAC per channel in month 3	Payback doubles and there is no scale
Churn above 5% a month	Cohort of 1,000 users over 6 months	**In Brazil LTV/CAC drops below 3x and the model stops closing**
The user will not let the agent pay	Measure authorisations accepted from week 1	The product stays a recommender and becomes a commodity
The 5% take rate in Brazil does not hold	Real pricing in 3 pilots with actual charging	Rail revenue falls by a third
The licence costs more or takes longer	Legal quote closed in both countries	More time in the partner model: lower margin

The scenario we have to face

In the **pessimistic scenario** (churn 6% / 5%, CAC 90/160, ARPU close to a third lower): the deficit at month 14 rises to **USD 553,269**, which this round still covers, but **there is no traction to raise Round 2: without it, cash runs out in month 16.**

Look at the detail: **in the bad scenario the problem is not the money, it is that the product does not catch on.** That is why the cut criteria are two, and both are early.

The ask

USD 700,000

for **14 months, up to Round 2**
Instrument and valuation: to be agreed with the investor
Target close: to be agreed

The two figures that frame everything

- **87% of the plan is building and complying, not growing.** Everything except acquisition: USD 591,353. Acquisition is USD 85,000, and it is minimal on purpose: **first we have to be able to charge a real user.**
- **The cash deficit at month 14 is USD 631,359.** This round covers it and leaves **USD 68,641** of buffer. **The amount comes from the spending plan, not from ambition.**

What the investor gets

→ Product executing real payments in the US and Brazil · **~459 users paying by month 12** · revenue USD 7,431/month · **CAC and churn measured with our own data** · 2 signed partners · written legal opinion.

USE OF FUNDS · 14-MONTH SPENDING PLAN (USD 676,353)



CLOSING

The vision: the agent as the gateway to the digital economy

Today

The agent pays for and manages the digital services of a person or a small business, in the US and Brazil.

Later

The same agent manages **goals**, not just orders — "reduce my digital spending this year", "renew what I use and cut what I do not" — and adds categories: transfers, insurance, investments and credit, always through authorised structures.

Further out

The engine (**agent + permissions + rail + history**) replicates to other corridors and other geographies. US–Brazil is the largest and hardest corridor: what works here works anywhere.

The user does not need to do more. They need to do less.

"You say what you need. The agent takes care of the rest."

APPENDIX

Agent architecture and permission layer

Components

- App and web · identity and KYB/KYC engine
- **Permission engine:** limits by category, amount and merchant; confirmation above threshold; revocability
- **Agent:** need → search → comparison → recommendation → authorisation → execution
- **Structured catalogue:** price, plan, currency, terms, availability by country and **actual successful payment rate**
- Ledger with balances per currency · payments engine and rail router
- **Risk and anti-fraud engine** · subscriptions and renewals engine
- **History and audit:** what the agent did, when, how much and under what authorisation

Build vs. buy

- **In-house:** agent, permissions, catalogue, experience, rules engine, history.
- **Third parties:** custody, initial licences, KYC, rails, issuing.

The asset that compounds

The catalogue is not a price list — that can be copied with scraping. It is **which services can be paid for from each country, at what real cost and with what success rate**: and only the one who executes the payments produces that.

APPENDIX

Sensitivity

CAC PAYBACK (MONTHS) · CONTRIBUTION MARGIN 72%

ARPU ↓ / CAC →	40	60	80	110	150
USD 8	6.9	10.4	13.9	**19.1**	**26.0**
USD 12	4.6	6.9	9.3	12.7	**17.4**
USD 15.20 · Brazil	3.7	**5.5**	7.3	10.1	**13.7**
USD 19.00 · US	2.9	4.4	5.8	**8.0**	11.0
USD 26	2.1	3.2	4.3	5.9	8.0

LTV / CAC · REFERENCE CAC USD 75

ARPU ↓ / Churn →	6%	5%	4.5%	3.5%	2.5%
USD 12	**1.9**	**2.3**	2.6	3.3	4.6
USD 15.20 · Brazil	**2.4**	**2.9**	**3.2**	**4.2**	**5.8**
USD 19.00 · US	3.0	3.6	**4.1**	**5.2**	**7.3**
USD 26	4.2	5.0	**5.7**	**7.3**	**10.2**

Read: the business lives or dies on **churn and CAC**, not on ARPU. One point of churn moves LTV/CAC more than any pricing lever.

APPENDIX

The six hard questions

1 · Why doesn't ChatGPT do it?

Because it does not pay or cancel. A model leaves you the link; we execute with your money and we answer for the result. On top of that, 62% of our revenue does not depend on recommending.

2 · You are in Chile: why the US–Brazil corridor?

Because that is where the pain is and that is where the money is: in Brazil payments are expensive and often the international payment fails, and in the US nobody knows what the company is paying for. What we bring is commercial execution and product, not physical presence: the platform is software, the partners are local and the team is hired where the volume is. Starting with the largest and hardest corridor is what makes the engine work afterwards anywhere else.

3 · How do you avoid the conflict of interest if you charge a commission?

We do not take a commission from the provider. We are paid through the rail and the savings: if I recommend badly, you do not save and I do not get paid.

4 · What happens if the take rate in Brazil falls to 2%?

Revenue from the rail falls by close to 60% and break-even moves further away: that is why the rail is not the only line. Savings and subscription do not depend on the take rate.

5 · Why two countries at once?

Because they are two different pains with the same product: in Brazil the problem is **being able to** pay; in the US the problem is **knowing** what to pay. The engine is the same.

6 · What stops a bank from copying it?

Because the bank competes for the transaction, not for the decision. Copying the agent requires permissions, catalogue and history: it is another product, with another team and another incentive.