Veterinary Entrepreneurship

Course in a Box: syllabus, instructor's guide, capstone, and AI track — to teach the method of The Entrepreneurial Veterinarian at any school in the world

PhD. Bernardo José Cortina Ceballos
DVM Sergio Armando Cortina Egea

Veracruz, June 2026

Serie Emprendedor · v1 · Base text: The Entrepreneurial Veterinarian (11 chapters)

Contents

Veterinary Entrepreneurship · Course in a Box

1 How to use this course in a box

This document is a **ready-to-teach instructor's kit**: everything you need to deliver the method of *The Entrepreneurial Veterinarian* as a course, without redesigning your own. It ships with a single underlying idea —**to drive the friction of adopting it down to almost zero**—, which is why it includes the syllabus, the session-by-session guide, the capstone project with its rubric, the breakdown of the AI track, and the map of alignment with the competencies your school must already meet.

It is **modular**: the same core assembles into two sizes —an **embedded module** of 6 weeks inside a related course (such as Livestock Production Systems) or a **full semester elective**—. Choose the assembly and follow its calendar.

> ✓ **Three rules to make it work**
>
> 1. **It is a course of doing, not of listening.** The axis is a real project; minimize lecture and maximize fieldwork and the workshop.
> 2. **Wide tent.** The book's examples are livestock-focused because that is where its author's feet are planted, but the method does not care about species: encourage every student to apply it to **their** area —small animal practice, pharmaceuticals, laboratory, aquaculture, beekeeping—. Adapt the local examples to your region.
> 3. **Honesty over the sale.** The method rests on validating with real data and conceding what must be conceded. Model that honesty; it is what sets a founder apart from a snake-oil salesperson.

The **base text** is the book (the 11 chapters) and the open resources at `serieemprendedor.com/recursos` (the validation canvas, the five-step checklist, and the linear-vs-exponential visualizer).

2 The course at a glance

2.1 Identity and promise

Field	Definition
Name (elective)	«Veterinary Entrepreneurship — "From specialist to founder: veterinary startups in the age of AI"»
Level	Final undergraduate semesters (or graduates, in a diploma format)
Language	Spanish (FMVZ-UV pilot); structure cloneable to English
Twofold promise	On completion, the student knows how to **validate a startup** and **use AI with judgment** — the two most in-demand competencies, through the veterinary lens
Format	Project-based (capstone), with fieldwork and AI workshops

2.2 Learning outcomes

On completing the course, the student is able to:

#	Learning outcome	Method phase
LO1	Distinguish a linear business (price-taker) from an exponential one and detect value-capture opportunities in their area	Red line
LO2	Validate an idea with real customers at low cost (Mom Test, Wizard of Oz, for under $50)	VALIDATE
LO3	Design an AI-native company (closed loop, team of agents) and translate their clinical judgment into instructions for the AI	BUILD
LO4	Build a defensible moat (data, network, trust, physical interface)	CAPTURE
LO5	Choose a business model and a financing path (OPC/ niche/exponential; SAS/SAPI/equity)	SCALE
LO6	Use AI tools with judgment: prompting, agents, limits, ethics, and human validation	AI track

LO7	Anticipate honest objections and make the decision to launch	Objections + Call

The alignment of each outcome with the formal accreditation competencies (AAVMC-CBVE, WOAH Day-1) is in Appendix 6.1.

2.3 The two assemblies

Assembly A — Embedded module (≈6 weeks, 12–18 h of contact time). To insert it inside a related course. It covers U0–U3 (in compressed form) and closes with a **mini-capstone**: validate **one** idea in a lightweight way. It is the «tasting» version —it plants the mindset and leaves a real deliverable—.

Week	Content	Deliverable
1	U0 Kickoff (linear vs exponential) + AI Level 0	3 candidate ideas
2–3	U1 SEE + U2 VALIDATE (Mom Test)	Decisive advantage + fieldwork (5+ conversations)
4	U3 BUILD (AI-native company) + AI lab	Sketch of the closed loop
5	Wizard of Oz + teamwork	Result of a Wizard of Oz test
6	Mini-pitch	3-min pitch + canvas

Assembly B — Full elective (15–16 weeks, 3–4 credits). The 9 units + full capstone + AI labs + a guest founder (or the field visits) + Pitch Day.

Week	Unit / activity	Capstone deliverable
1	U0 + team formation + AI Level 0	Team + 3 candidate ideas
2–3	U1 SEE	Chosen idea + **decisive advantage** + canvas v1
4–5	U2 VALIDATE (fieldwork)	Evidence of 10+ conversations (Mom Test)
6	AI lab + Wizard of Oz	Result of a Wizard of Oz test
7–8	U3 BUILD	AI-native design: closed loop + map of agents
9	U4 CAPTURE	Moat definition
10–11	U5 SCALE	Business model + financing path
12	U6 Cases and market	Tomography of a real case
13	U7 Objections (objection clinic)	Response to your 3 toughest objections
14	Pitch rehearsal + peer feedback	Pitch v1

15–16	U8 The call + **Pitch Day** (external judges)	**Final pitch (5 min) + validated canvas**

2.4 Assessment

Component	Weight
Capstone (deliverables by phase + final pitch)	60%
AI labs + portfolio	20%
Participation (objection clinic, peer feedback, fieldwork)	20%

No traditional exam: the axis is the project. **«Mixed» variant** if your school requires an exam: 50% capstone + 20% labs + 20% participation + **10% concept quiz** (market, AI-native company, business models).

3 The nine units

Each unit states the book chapters it covers, its objective, the key content, the in-class activity, the stretch of the **AI track** it draws on, the capstone deliverable, and a **note for the instructor**.

3.1 U0 · Kickoff: linear vs exponential

Chapters 1–2. Objective: for the student to understand that the difference between linear and exponential is not one of size but of **nature**, and why ceasing to be a price-taker (of a *commodity*) is the only red line. **Key content:** your knowledge as a convertible asset; the ceiling is the vision, not the resources; the linear-vs-exponential chart; the red line of the *commodity*. **Activity:** the linear-vs-exponential visualizer on the landing page; each student classifies their «dream job» as linear or exponential. **AI track — Level 0 (kickoff):** first exposure —what generative AI is and is not—. **Deliverable:** team formed + 3 candidate ideas.

> ✓ **For the instructor**
>
> Open with a real story that stretches the tent: the colleague from the book's first page, a bee veterinarian who applied the method outside livestock and became a national authority. The day-one message is «this is for you, whatever your species.»

3.2 U1 · SEE: the exponential opportunity

Chapter 3. Objective: detect a value-capture opportunity in their area and formulate their **decisive advantage**. **Key content:** the Antimicrobial Resistance Network («the Waze of resistance») as a truly exponential example; the network effect (Metcalfe's Law, value grows with the square of the users); the tool that is useful «single-player» before the network exists (the cold start); privacy by aggregation; your clinical eye as the scarce data. **Activity:** with the validation canvas, each team sketches its idea and its decisive advantage. **AI track:** seeing with AI —analyzing a dataset from the ranch or clinic to find a pattern—. **Deliverable:** chosen idea + decisive advantage + canvas v1.

> ✓ **For the instructor**

Insist that the student's example **does not have to be livestock**. A good decisive advantage is born from something the student already knows or holds and an outsider does not: access to a guild, a clinical dataset, the trust of a region.

3.3 U2 · VALIDATE: real customers, low cost

Chapter 4. Objective: validate the idea with real customers before investing. **Key content:** the *Mom Test* (how to ask without prompting the answer you want to hear); the *Wizard of Oz* (simulating the product by hand before building it); validating for under $50; the lighthouse ranch; «boots in the mud, not mass email.» **Activity:** fieldwork —each team holds 10 or more real conversations with potential customers—. **AI track:** validating with AI —generating the interview script and synthesizing the Mom Test findings—. **Deliverable:** evidence of 10+ conversations.

✓ **For the instructor**

The number-one trap is asking leading questions («wouldn't an app that… be great?»). Teach the Mom Test **before** the fieldwork and have each team show you their script. The goal is not for them to be told yes; it is to learn the truth cheaply.

3.4 U3 · BUILD: the AI-native company

Chapter 5. Objective: design a company that runs on AI. **Key content:** AI as an **operating system**, not as a copilot; «from your intuition to an instruction» (turning a clinical decision into an *if… then* diagram, which is exactly what the agent reads); the *Planner–Critic–Executor* team of agents; the company as a **closed loop** that self-regulates; «count tokens, not payrolls.» **Activity / AI lab:** prototype an agent flow with no-code tools (e.g., n8n) for a capstone task. **AI track — Level 1 (core):** build with agents. **Deliverable:** AI-native design —the closed loop and the map of agents—.

✓ **For the instructor**

The concept that is hardest and worth the most is the **Critic**: the agent (and, above all, the human) that audits the work before it goes out. It is the *human-above--the-loop* —human validation above the loop— and it is the heart of the method's technical honesty. Spend time on it.

3.5 U4 · CAPTURE: the defensible moat

Chapter 6. Objective: build an advantage that cannot be copied. **Key content:** the four--layer moat —the accumulated data, the network effect, the trust of the guild, and the physical interface between the *bit* and the atom—; why an engineer with more capital cannot steal it from you. **Activity:** each team defines its moat and subjects it to the «giant test» —would a corporation with better AI replicate it overnight?—. **AI track:** data as a moat; how it accumulates and is protected. **Deliverable:** moat definition.

> ✓ **For the instructor**
>
> Distinguish a **moat** from a **temporary advantage**. «I got there first» or «my app is prettier» are not moats; a collective dataset that grows with every use is. If the team cannot name why its advantage **deepens** over time, it does not yet have a moat.

3.6 U5 · SCALE: model and financing

Chapter 7. Objective: choose the business model and the financing path. **Key content:** choose your ceiling —a one-person company (OPC), a profitable niche, or exponential— but with your north star on the exponential; the legal tools by stage (SAS, SAPI); venture capital and **the VC's noose** (dilution, control, pressure to grow); «Delaware is not the escape: it is the funnel.» **Activity:** each team chooses its model and traces its financing path. **AI track:** «count tokens, not payrolls»; open-weight models so as not to depend on a single provider. **Deliverable:** business model + financing path.

> ✓ **For the instructor**
>
> Be honest about the noose of venture capital: the rocket has a price. And make it clear that choosing a profitable niche **with a moat** is not settling —it is ownership —; what the method fights is remaining a price-taker, not the size.

3.7 U6 · Cases and market

Chapters 8–9. Objective: analyze real cases with method and locate where to enter. **Key content:** five startup cases with the **Tomography Method**; the global market map; the six--filter framework for choosing a beachhead. **Activity:** each team performs the «tomography»

of a real case from their area. **AI track:** analyzing a case with AI assistance. **Deliverable:** tomography of a case.

> ✓ **For the instructor**
>
> Bring cases from your students' region, not just the book's. A local case —successful or failed— is worth ten distant examples, and it reinforces the wide tent.

3.8 U7 · Honest objections

Chapter 10. Objective: anticipate and answer the objections, and face obsolescence head--on. **Key content:** the eleven cracks in the project (from «ranchers don't adopt technology» to «what if the AI learns your very own clinical eye?»); obsolescence as an honest hypothesis —if AI keeps up this pace, dig deep moats so that irrelevance takes its time—. **Activity:** the **objection clinic** —each team faces the three toughest objections to its project; classmates fire them at them—. **AI track:** ethics, limits, and the future of work. **Deliverable:** a written response to their three toughest objections.

> ✓ **For the instructor**
>
> Model honesty: conceding what must be conceded **strengthens** the proposal, it does not weaken it. A founder who knows the cracks in their own project inspires more confidence than one who swears everything will turn out fine.

3.9 U8 · The call and Pitch Day

Chapter 11. Objective: close with the decision to launch and communicate it. **Key content:** the Monday Protocol (the first concrete move); the student's knowledge as the scarce input; the `serieemprendedor.com` community. **Activity: Pitch Day** —a 5-minute presentation before external judges—. **Deliverable:** final pitch + validated canvas.

> ✓ **For the instructor**
>
> Invite real and diverse judges —an entrepreneur, someone who understands capital, and a guild customer (a rancher, a clinic owner)—. Let the last experience of the course be defending a **validated** idea before the real world, not before the instructor.

4 The capstone project

The capstone is the axis of the course: teams of 2 to 3 students validate a **real idea from their own area** by running it through the method. **It is not a desk business plan:** it demands getting out into the field and talking with real customers. The final product is not «the best idea,» but **the best validated learning**.

4.1 Deliverables by phase

Phase	Deliverable
SEE	Chosen idea + **decisive advantage** + validation canvas v1
VALIDATE	Evidence of 10+ real conversations (Mom Test) + result of a Wizard of Oz test
BUILD	AI-native design: the closed loop and the map of agents
CAPTURE	Moat definition (and why it deepens over time)
SCALE	Business model + financing path
Closing	**Final 5-min pitch + validated canvas + validation evidence**

4.2 Capstone rubric

Five equally weighted criteria. Each is scored 1–4 (insufficient, developing, competent, outstanding).

Criterion	What is assessed
1 · Problem and decisive advantage	Is the problem real and is the decisive advantage something an outsider does not have?
2 · Rigor of validation	Did they talk with real customers, without leading questions (Mom Test)? Does the evidence rule over opinion?
3 · AI-native design and moat	Do the closed loop and the team of agents make sense? Does the moat deepen with use?
4 · Honesty	Do they acknowledge their assumptions and answer their toughest objections without window-dressing?
5 · Pitch communication	Is the pitch clear, honest, and memorable in 5 minutes?

★ **Golden Rule: The criterion that weighs most in practice**

Criterion 2 —**rigor of validation**— is what separates an entrepreneurship course from an ideas contest. It rewards the team that discovered its idea was wrong and **pivoted with evidence** over the one that defends a pretty idea without having talked to a single customer. Failing cheaply and learning is winning.

5 The two-level AI track

The course teaches how to harness AI at **two levels at once**: it lifts off the floor anyone who has never used it and, at the same time, teaches those who already use it how to apply it to their startup. It is the differentiator that no traditional «practice management» class offers.

5.1 Level 0 — Fluency (for those who have never used it)

It is concentrated at the start (U0) and remains as reference material for the whole course. It covers, without button names that will go stale:

- **What generative AI is and is not.**
- **Prompting as delegation:** speaking to AI as you would to a brilliant but literal-minded intern.
- **Search engine vs. copilot vs. agent:** three distinct ways of using it.
- **Where it fails** (hallucination) and why **you** always validate.
- **Your data and privacy:** what not to upload, how to aggregate.
- **Open vs. closed models** and why they matter to the founder.

- **Subtle point: One source, three uses**

This Level 0 is exactly the «AI essentials» that the book should carry as a **perennial appendix**: the same content serves the book, this course, and the online resources —and, being conceptual rather than tool-based, it does not go stale with each new version of the apps—.

5.2 Level 1 — AI applied to the method (for everyone)

Woven into each phase of the method; each stretch is practiced in a lab:

Phase	AI applied
SEE	Analyze ranch or clinic data to find patterns and opportunities
VALIDATE	Generate interview scripts; synthesize the Mom Test; design a «human + AI» Wizard of Oz
BUILD	Design the closed loop and the team of agents (Planner–Critic–Executor); no-code; an MVP

CAPTURE	Set up the flow that accumulates the data that becomes a moat
SCALE	Estimate the cost in tokens; use open models so as not to depend on a provider

5.3 The labs

Throughout the course, 3–4 **hands-on labs** in which the student **actually uses** an AI for a task of the method: building their interview script (U2), prototyping an agent flow (U3), setting up a mini data pipeline (U4). Each lab produces an artifact that goes into the portfolio.

> ▲ **The golden rule of the AI track**
>
> The AI **executes**; the veterinarian **cures and validates**. Every AI output passes through human judgment before it touches an animal or a customer. Teaching the tool without teaching this limit would be irresponsible —and would break the honesty that is the method's hallmark—.

6 Appendices

6.1 Appendix A — Map of alignment with competencies

This is the course-level version of the **alignment map** —the highest-leverage artifact for curricular adoption—. It shows how each learning outcome feeds a competency that the accreditation frameworks **already require**. (The fine-grained map —chapter by chapter against numbered competencies— is produced separately for each school's curriculum committee.)

LO	Learning outcome	Formal competency it feeds
LO1	Linear vs exponential; value capture	CBVE Domain 8 «Financial & Practice Management»; WOAH adv. 3.8 «Administration and management»
LO2	Validate with real customers	CBVE D8 «applies financial principles to decisions»; D7 «Professionalism»
LO3	AI-native company	CBVE D8 (management/innovation) + cross-cutting AI competency
LO4	Defensible moat	CBVE D8 (business sustainability)
LO5	Model and financing	CBVE D8 (revenue, expenses, compensation); RCVS «Business/Finance»
LO6	AI with judgment	Cross-cutting AI competency; WOAH «communication»
LO7	Objections and the decision to launch	CBVE D7 «career planning» with financial implications

- **Why this appendix matters**

The world's four major frameworks (AAVMC-CBVE in the U.S./Canada, RCVS in the United Kingdom, AVBC in Australasia, WOAH as the global standard) **already require** a domain of business and finance. This course does not ask the school to add a new topic: it offers **the resource that fills a domain it already has to meet** —and it takes it from the mandatory minimum (financial literacy) to the exponential —. That is the honest framing and, at the same time, the most persuasive one before a curriculum committee.

6.2 Appendix B — Materials and resources

Resource	Use
The book *The Entrepreneurial Veterinarian* (11 chapters)	Base text; one unit per method block
Validation canvas (`serieemprendedor.com/recursos`)	Axis of the capstone, from U1
Five-step checklist	Guide for the VALIDATE phase
Linear-vs-exponential visualizer	U0 activity
90-min class script + deck (15 slides)	Basis of the opening session / express version
Community `serieemprendedor.com/comunidad`	Closing: from the classroom to the founder network

6.3 Glossary of the method

Decisive advantage — what you have or know and an outsider does not (formerly «unfair advantage»). · **Closed loop** — a company that measures, decides, and acts in a self-regulated way, like homeostasis. · **Team of agents (Planner–Critic–Executor)** — the three-role AI pattern: plans, executes, and audits. · **Human-above-the-loop** — the human judgment that validates above the automatic loop. · **Mom Test** — a technique to validate without asking questions that prompt the desired answer. · **Wizard of Oz** — simulating the product by hand before building it. · **Moat** — an advantage that deepens over time (data, network, trust, the physical). · **Beachhead** — the first niche through which one enters the market. · **Tomography Method** — a methodology for analyzing a startup in layers. · **OPC** — a one-person company. · **Price-taker** — one who does not set their own price (the red line of the *commodity*). · **CBVE** — *Competency-Based Veterinary Education*, the AAVMC competency framework.