ADSA Annual Meeting • Production, Management & the Environment

Optimizing Dairy Cow Replacement Decisions Using Deep Reinforcement Learning

A feasibility study in an evolving, stochastic dairy environment

Yijing Gong, Lazaro Henrique da Silva, Victor E. Cabrera

Department of Animal & Dairy Sciences, University of Wisconsin–Madison

WISCONSIN
UNIVERSITY OF WISCONSIN–MADISON

Replacement: a cost and a herd-improvement lever

(USDA, 2016)

~35%

cows replaced per year (US)

(Hawkins et al., 2020)

2nd

largest operating expense on a dairy, after feed

(Buterbaugh et al., 2026)

~2.9

Average lactations before a cow leaves (US)

(Schmitz, 2025)

>$3,000

recent replacement heifer price

The opportunity

Every cull is a chance to upgrade genetics, health & productivity.

BACKGROUND

Two approaches for the same decision

Dynamic Programming (DP)

- The established approach for ~40 years.
- Built from a full specification of the system's transition probabilities.
- Computes the exact optimal policy.
- Well understood — a strong benchmark.

Deep Reinforcement Learning (DRL)

- A learning-based approach.
- Learns a policy by interacting with the environment.
- A neural network generalizes across similar cow states.
- Established in other fields; little explored in dairy.

Our stance: DP is the benchmark. We test DRL as a complementary approach worth exploring — not a replacement.

OBJECTIVE

A feasibility study — three questions

A proof of concept

Learn

Can a DRL agent learn profitable, biologically rational keep / replace policies?

Generalize

Do those policies hold up across five distinct market scenarios?

Adapt

Can it adjust after an economic shock — without re-specifying the model?

METHODS

Cow replacement as a Markov decision process and reinforcement learning (RL) loop

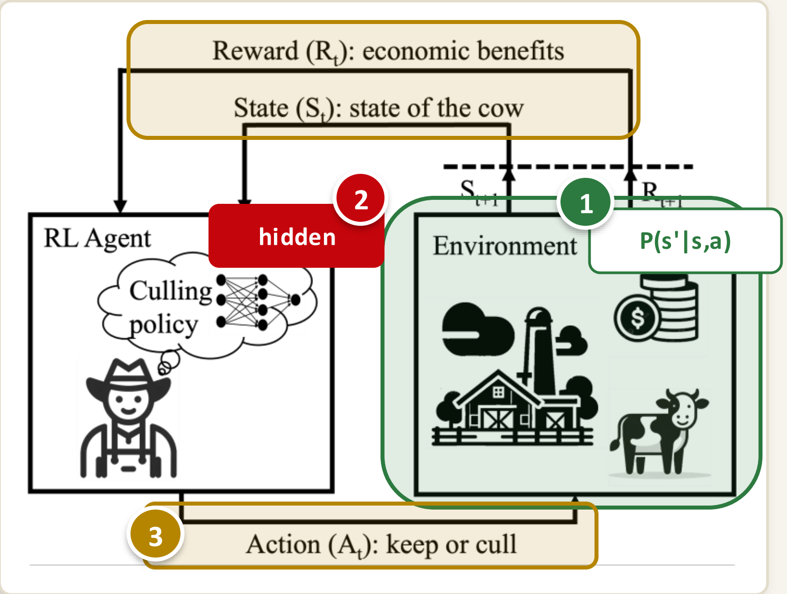

1 P(s'|s,a) builds the environment, shaping the cow's life trajectory.

2 It stays hidden from the agent.

3 The agent learns by interacting, not from the model.

State — the cow

Parity (1–12) · Month after calving (1–20)
Month in pregnancy (0–9) · Clinical mastitis (0/1)

Action

Keep, or replace with a springing heifer

Reward (monthly partial economic return)

Keep: milk + calf − feed − breeding − treatment
Replace: slaughter − heifer cost − heifer feed

Objective: Maximize the expected discounted sum of economic returns over the planning horizon

METHODS

Deep Q-Network agent & five economic scenarios

- Deep Q-Network (**DQN**): approximates the action-value function, Q(state, action), using a neural network.
- 500,000 training episodes per scenario · 180-month horizon · 5 random seeds.

Heifer supply × milk market

+ milk price
- feed price

+ milk price
- feed price

	Good milk market	Bad milk market
Heifer oversupply	**OG**	**OB**
Heifer undersupply	**UG**	**UB**

- replacement heifer cost
- calf price
- slaughter price

+ replacement heifer cost
+ calf price
+ slaughter price

+ BL (2025 baseline market conditions)

Stable learning with market-sensitive profitability

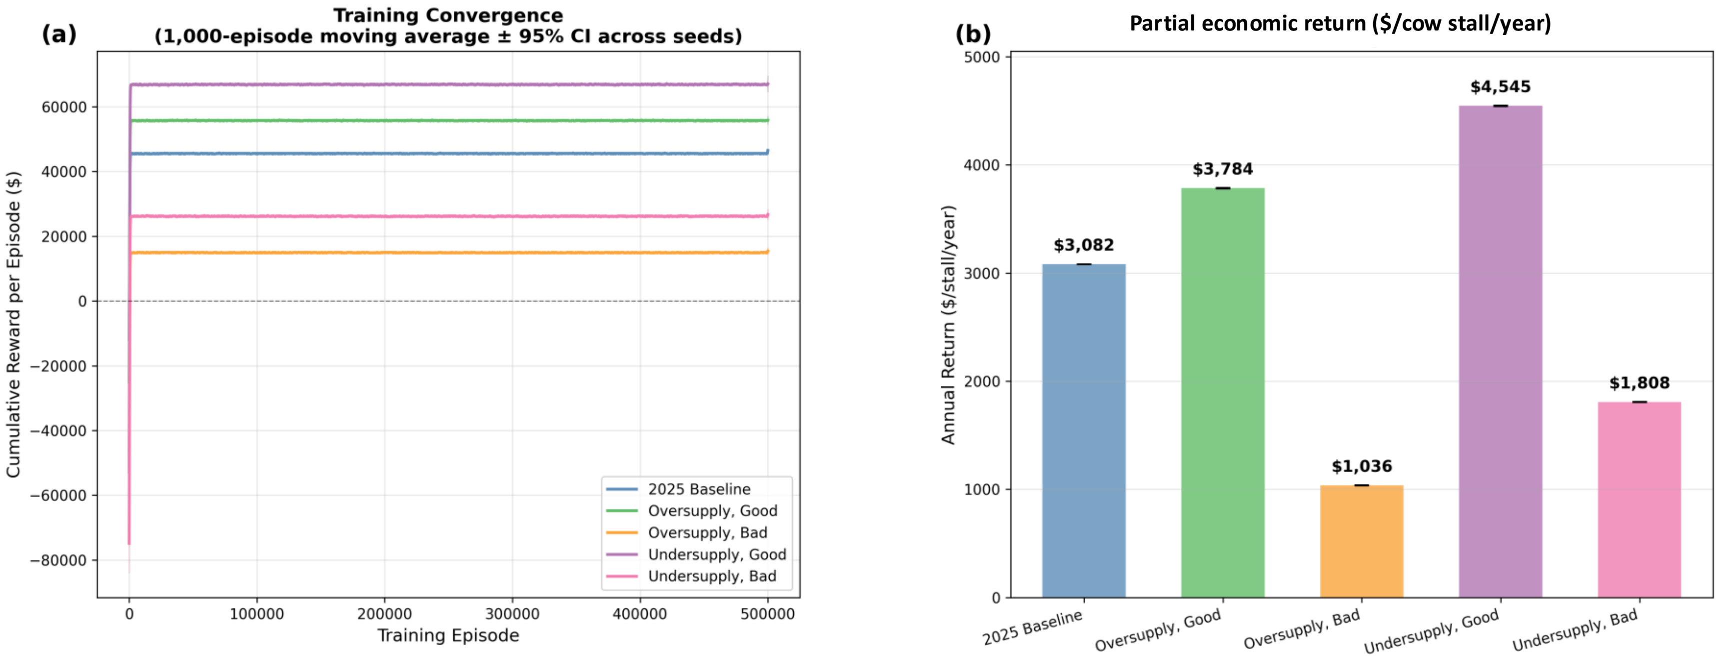

Market conditions shape culling and herd dynamics

Scenario	IOFC ($/stall/yr)	Total culling (%)	Mean parity	Productive life (yr)
BL — 2025 baseline	**3,082**	**28.5**	**2.71**	**3.5**
OG — over · good milk	3,784	21.4	3.40	4.7
OB — over · bad milk	1,036	21.1	3.43	4.7
UG — under · good milk	4,545	22.8	3.27	4.4
UB — under · bad milk	1,808	23.9	3.08	4.2

Culling does not respond linearly to market conditions:

- Baseline (BL) culls the hardest and maintains the youngest herd, despite intermediate economic values
- Heifer oversupply scenarios (OG and OB) cull the least, likely because lower salvage and calf values

RESULTS · 3

The learned policy is biologically rational

Market condition: Baseline
Green = keep; **Red** = replace

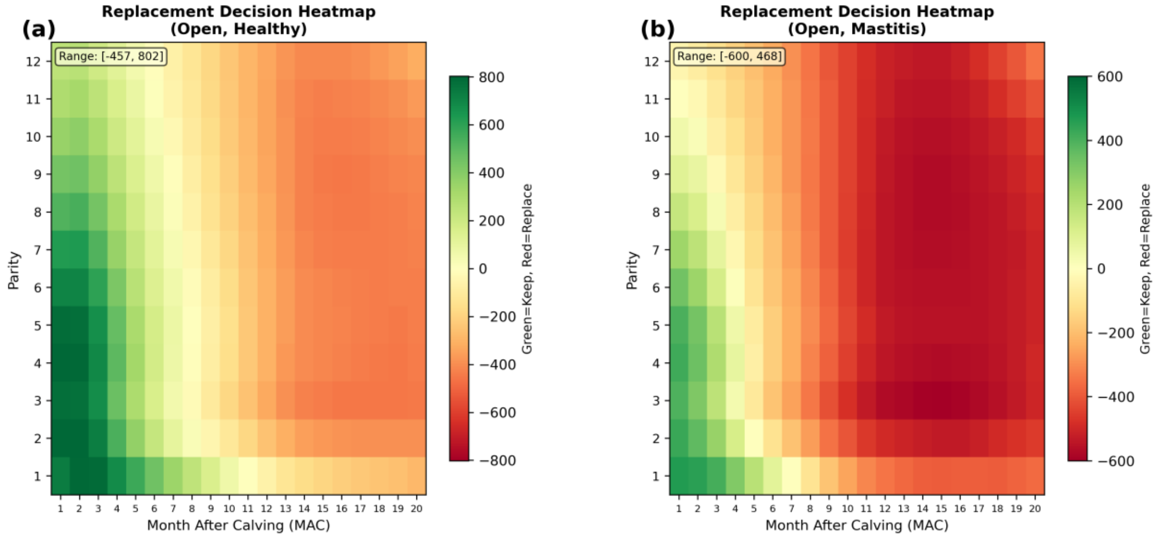

- **Younger** cows → retain
- **Yellow** → decision flexibility
- **Clinical mastitis** → cull ~3–4 months earlier

Putting a dollar value on biology

Value of a new pregnancy (BL)

- Healthy cows averaged \$418 (parity 1), \$313 (parity 2), \$231 (parity 3).
- Pregnancy raises retention value; open cows lose value as lactation drags on.
- Timing & parity both matter — a single number understates it.

Cost of clinical mastitis

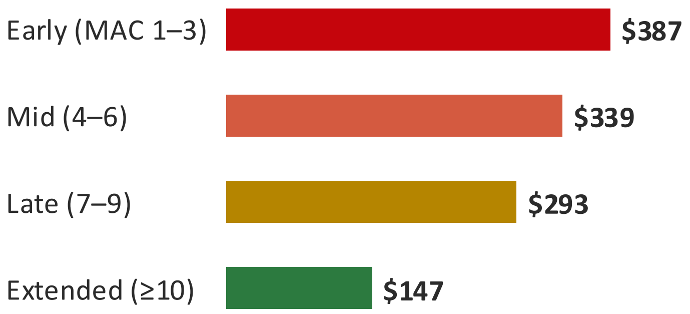

Highest in middle-parity cows.

Adapting to an unforeseen economic shock

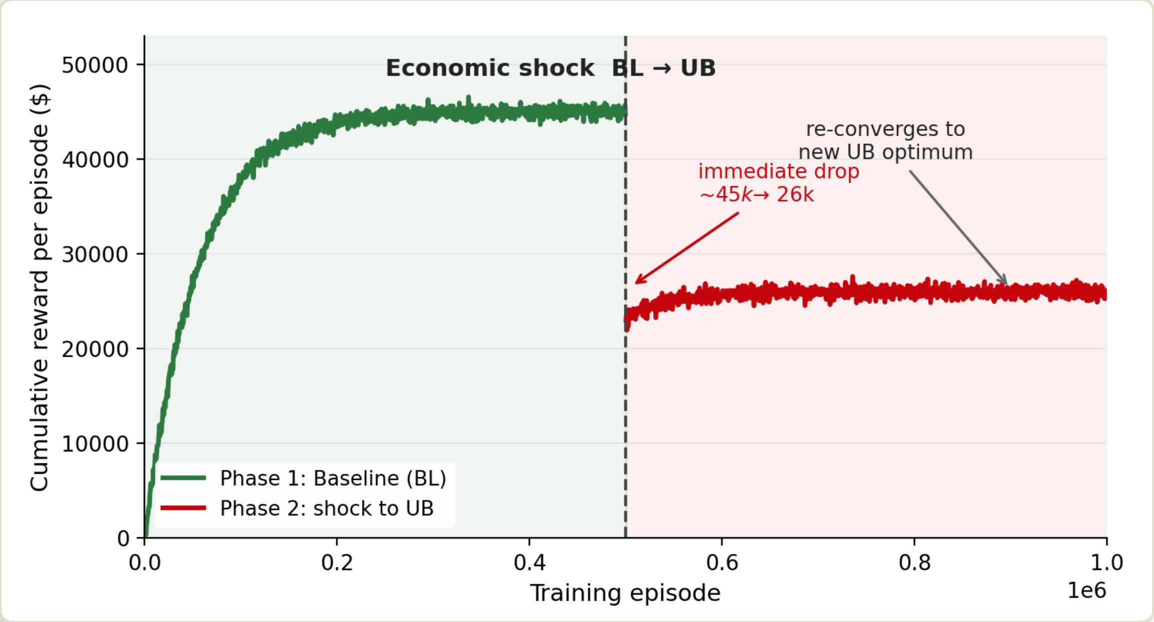

- Agent trained under BL, then the market flips to UB mid-training.
- Reward drops immediately (~$45k → ~$26k per episode).
- It keeps learning and re-converges to the UB optimum.
- Final herd & policy match an agent trained on UB from scratch.

Scope of this study: what we claim and what we do not

What we show

- DRL learns profitable and sensible culling policies
- Recommendations shift with market, pregnancy, and disease status
- The model adapts online to economic shocks

Limitations → next steps

- DP benchmark with identical parameters underway
- Only clinical mastitis; no lameness or other disorders.
- No within-lactation yield variation.
- Next: learn from farm data
    - No predefined transition probabilities
    - Higher-dimensional decisions
    - Continuous model updating

TAKEAWAYS

Deep RL is a viable foundation for dairy replacement decisions

- A DQN agent learned profitable, biologically rational keep/replace policies across five market scenarios.
- It quantified the value of pregnancy and the stage-dependent cost of mastitis and adapted online to an economic shock.
- A step toward data-driven, adaptive decision-support tools.

Victor E. Cabrera vcabrera@wisc.edu • University of Wisconsin–Madison


EXHIBIT HALL
200 LEVEL MEETING ROOMS
HILTON SKYWALK
NORTH BLDG
WISCONSIN AVE. EXIT
SPANISH
PORTUGUESE
MANDARIN