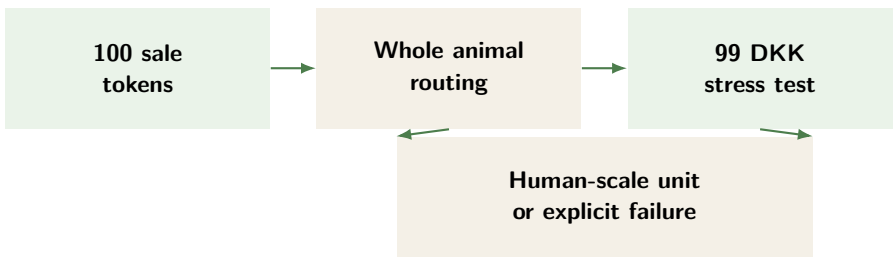


100 Burger a Day

Designing a human-scale burger business around the whole animal

Can a burger business deliberately stop at 100 sales, pay human work and suppliers properly, and still make economic sense?



Author Dennis Hedegreen
Publisher Hedegreen Research
Series Market Gaps
Evidence M2 / P0

Status Release candidate
Date 10 August 2026
Version RC1
DOI Not assigned

Abstract

Burger businesses are normally discussed through throughput, menu price, food cost and location. This paper begins with two deliberately chosen constraints: a reference unit may sell no more than **100 premium combos per opening day**, and the complete combo is tested against a **99 DKK including VAT** price target. Neither number is claimed to be optimal. They are design constraints used to force the model to confront bounded human work, whole-animal procurement and a normal consumer-facing price.

The reference design uses 100 premium sale tokens, a 140 g raw beef patty, a simple burger-fries-house-soda menu, and explicit valuation of owner/operator labour. Instead of treating ground beef as an isolated ingredient, the model begins with a whole reference animal and routes premium burger material, protected cuts, secondary beef products and rendered fat separately. All carcass figures remain provisional until observed breakdown data exist.

Under the current routing assumptions, the model yields **~412 nominal patties per reference animal** and uses a provisional reserve/variance allowance to plan on **~400 saleable burgers**. Dry-run simulations indicate that a two-operator service line is more robust than a one-person line under a short 50-combo-per-hour rush. At five sold-out opening days per week, upstream demand also becomes unusually legible: 500 buns and 75 kg of served fries, with similarly bounded requirements for other inputs.

The model does not yet prove commercial viability. The 99 DKK case depends on real ingredient invoices, fixed overhead, average utilization and the realized value of secondary animal products. The farmer-benefit thesis is also conditional: a higher gross carcass payment matters only if additional obligations do not erase the farmer's net gain. Physical throughput, customer demand and coordination burden remain empirical-open.

The contribution is therefore not a restaurant success claim but a reproducible design method: chosen constraints, explicit accounting boundaries, falsification rules, implementation artifacts and a physical pilot protocol are exposed so another operator can test the system without requiring permission from its author.

EVIDENCE STATUS

Release Candidate 01 is an M2/P0 implementation-ready design, not an operating case study. The unified economic ledger is included. K1 remains OPEN/CONDITIONAL; K2 and K3 have quantitative floors but remain empirically open. No restaurant, farmer contract, observed utilization curve or replicated unit is claimed.

The Model in Two Minutes

What is the idea?

A small burger unit deliberately exposes only **100 premium sale tokens per opening day**. The reference combo is a burger, tallow fries and house soda tested at **99 DKK including VAT**.

Why stop at 100?

Because the cap makes inventory, human work and supplier demand finite enough to inspect. The goal is not maximum throughput. It is a productive unit that can stop.

Why the whole animal?

A cow is not 400 burger patties. If the restaurant wants to pay the farmer more by using the animal better, the rest of the animal must have real channels and real value.

Can 99 DKK work?

Still plausible, but not yet passed. The earlier 70-sales/day claim mixed an opaque food-cost placeholder with a separate farmer-premium model. The unified ledger now makes viability depend jointly on utilization, non-beef cost and realized secondary contribution.

What is the hardest problem?

At five sold-out days, the model requires about **1.25 animals/week** and creates roughly **221 kg/week of secondary edible output** that must find credible buyers.

What would kill it?

Low utilization, unsold secondary products, no farmer net benefit, too much human work, poor quality at peak, or excessive coordination.

What are we giving away?

The model, pilot protocol, supplier interfaces, economic framework and failure tests are intended to be public enough for another operator to try without permission.

**The point is not to prove that a burger bar works on paper.
The point is to make a productive unit easy to understand, test and falsify.**

How to Read This Case

MARKET GAPS

Hedegreen Research uses Market Gaps for productive structures that appear missing, undersupplied or needlessly difficult to coordinate. A case publishes the model, economics, human-work boundary, interfaces and failure tests rather than treating the idea itself as scarce property.

- | | |
|------------------------------|---|
| 1. The gap | The market gap and the question being tested. |
| 2. Why 100? | The cap as inventory and labour architecture. |
| 3. Why the whole animal? | Carcass routing and the 413-to-400 planning bridge. |
| 4. Farmer + secondary burden | Farmer economics and the 221 kg/week bottleneck. |
| 5. Can 99 DKK work? | One ledger linking K1, K2 and K3. |
| 6. Human-scale operation | Human work, throughput and quality. |
| 7. Local nodes | Buns, potatoes and bounded supplier interfaces. |
| 8. Thin coordination | State-sharing without a headquarters owning trade. |
| 9. Ownership | Governance boundary. |
| 10. Kill conditions | Evidence classes and K1-K6. |
| 11. Try it | Public implementation and the next physical test. |

READER RULE

Understand the business question before the evidence machinery. Evidence labels and kill conditions make the intuition testable; they should not be required to discover what the idea is.

1. Why This Business Exists

The conventional question for a burger business is usually some version of: **how do we sell more burgers?** More footfall, more delivery orders, more opening hours, more menu extensions, more stores. 100 Burger a Day starts by refusing that optimization target. Its first deliberate constraint is that a reference unit exposes only 100 premium sale tokens on an opening day.

The second constraint is equally artificial by design: the complete combo is tested against a 99 DKK including-VAT price target. The project is not claiming that 100 is the profit-maximizing volume or that 99 DKK is the market-clearing price. The point is to ask whether a useful small productive unit can survive inside those boundaries.

That creates a different research problem. Burgers are standardized restaurant products, but cattle are not born as standardized patties. Beef production begins with an animal and passes through slaughter, chilling, fabrication and sale of multiple products with different values. A restaurant can avoid most of that complexity by buying exactly the mince or cuts it wants. That is operationally sensible. It also means the restaurant normally has little reason to ask whether a different use of the whole animal could support a better upstream payment.

There is a second mismatch: small owner-operated businesses can turn attractive gross margins into long working weeks. Eurostat's 2024 labour-force data reported that 27.5% of self-employed workers worked 49 hours or more per week, compared with 3.4% of employees.¹ This is not evidence about Danish burger operators specifically. It is enough to justify a stricter accounting rule: if a proposed small business works only because owner time is valued at zero, the economics are misleading.


The central research question is:

Can a capped, farmer-linked burger unit create enough value from the whole animal to pay suppliers fairly while remaining worthwhile for a small owner-operator team?

No 100BAD unit is operating. No farmer contract has been signed. No 100-order service day has been observed. No real utilization curve exists. The current maturity is therefore **M2/P0**: implementation-ready design, not operating evidence.

HOW THIS PAPER TESTS ITSELF

The design loop is: choose constraints -> model the productive unit -> expose hidden work -> attack the model -> define kill conditions -> run the smallest physical test -> revise or abandon. Quantitative claims are tagged OBSERVED, MODELED, EXTERNAL, SCENARIO, DESIGN, HYPOTHESIS or SPECULATIVE.



Evidence: DESIGN/A. M2 is not observed operation.

Figure 1: Evidence maturity. 100 Burger a Day is currently an implementation-ready design case (M2), not an operating or replicated case.

¹Eurostat, "27.5% of self-employed people worked 49 hours or more per week in 2024," 21 August 2025. Eurostat source

2. Why Stop at 100?

The reference unit exposes exactly **100 premium sale tokens per opening day**. A sale token is the shared inventory truth across every ordering surface. If five customers prepay for pickup, the count falls by five. If the counter says 37, the system must be capable of fulfilling 37 more premium orders.


Sale inventory is separated from recovery inventory. A kitchen may hold extra patties for a failed quality check, a remake, a pre-opening taste check or an explicitly coded staff meal. Those are not burgers 101-104. The workpack therefore tracks:

`reserve factor = premium patties consumed / premium burgers sold`

A reserve factor above 1.00 feeds back into daily food cost and annual animal demand. For planning only, this paper uses a **~3% provisional planning allowance** (SCENARIO/D) to bridge nominal carcass yield and saleable planning yield. It is not an observed kitchen loss rate; the exact reproducibility value is kept only in the technical workpack.

The cap exists for four reasons: inventory readability, bounded production demand, upstream planning and a clean stopping condition. It does **not** prove that the total workday ends when burger number 100 is sold. Prep, receiving, close, cleaning, admin, maintenance and exceptions still count.

For a fixed location, the preferred front-end hypothesis is a Japanese ramen-style ticket flow. The customer chooses, pays and receives a ticket or number. The kitchen receives a paid order and the sale-token counter decrements automatically. The point is not to eliminate choice but to move payment and order-entry away from the human production bottleneck.



Evidence: DESIGN/A. 100 sold != 100 patties consumed.

Figure 2: The daily inventory boundary. Premium sales draw from one shared pool of 100 sale tokens; recovery/QC inventory is separate.

3. Why Start With the Whole Animal?

The reference design uses a simple maxim:

The burger is the anchor product. The animal is the unit of production.

The current carcass model begins with a 600 kg live reference animal and uses conventional dressing and cooler-shrink assumptions. Instead of repeating low-confidence decimals in the main narrative, the publication-level interpretation is simpler: the model assumes a chilled carcass in the mid-300 kg range and an edible/freezer-ready output around **235 kg**.

Mississippi State University Extension provides a useful external sanity check, describing average beef dressing percentages around 60-64%, cooler shrink around 2-5%, and a rough rule of thumb of approximately 40% of

live weight becoming freezer-ready beef.² The 100BAD edible-pool assumption is close to that aggregate rule. This supports plausibility, not precision.

The premium burger uses a **140 g raw patty** as a reference specification. The current technical formulation is chuck-led with fat balancing while protecting selected higher-value muscles where separate sale is rational. Sensory questions - including the use of brisket fat - remain open.

Under current routing assumptions, the premium blend is about **57.75 kg**, which at 140 g yields **~413 nominal patties**. Applying the provisional **~3% planning allowance** gives **~400 saleable planning burgers per reference animal**. The first number is physical arithmetic on a modeled blend; the second is a conservative planning yield. Both remain low-confidence until real carcasses and reserve consumption are measured.

The remaining edible output is not credited automatically as value. It creates an economic burden that must be absorbed through plausible channels such as steaks/roasts, a standard frozen-patty stream, slow-cook cuts, tallow or other products.

The key rule added after adversarial review is:

No secondary-product value may be credited to the premium burger economics unless a plausible sales channel exists for that value.

This turns secondary absorption from an optimistic side assumption into a central empirical gate.

THE PLANNING BRIDGE

About 413 nominal 140 g patties come out of the current modeled premium blend. A **3% provisional planning allowance** converts that into about 400 saleable planning burgers per animal. It is a planning placeholder, not observed waste.

Table 1: Reference animal assumptions. Low-confidence values are deliberately rounded.

Variable	Publication value	Evidence	Basis
Live reference animal	600 kg	MODELED/D	Project reference
Dressing context	60–64%	EXTERNAL/B	MSU Extension
Edible/freezer-ready pool	235 kg	MODELED/D	Model + aggregate sanity check
Premium blend	58 kg	MODELED/D	Project routing
Premium patty	140 g raw	DESIGN/B	Product specification
Nominal premium output	413 patties/animal	MODELED/D	57.75 kg / 140 g
Planning allowance	3%	SCENARIO/D	Rounded publication parameter
Saleable planning yield	400 burgers/animal	MODELED/D	413 less provisional allowance

²Mississippi State University Extension, “How Much Meat to Expect from a Beef Animal: Farm-Direct Beef.” Used here only as an aggregate technical sanity check, not as observed Danish carcass data. MSU Extension source




Figure 3: Whole-animal routing hypothesis. Nominal physical yield and saleable planning yield are shown separately; secondary value is credited only when a plausible channel exists.

4. Can Paying the Farmer More Work?

The economic hypothesis is conditional:

If the system creates more realized value from the animal, part of that value may be available for a better farmer payment.

That is different from claiming that the farmer is automatically better off.

For a dated Danish benchmark, Danish Crown's week-33 2026 cattle settlement implies approximately **15.52k DKK** for the project's modeled reference carcass.³ The workpack tests a **+15% gross farmer-payment scenario**, which raises that figure to approximately **17.85k DKK**.

The difference is about **2,328 DKK per animal**. At five sold-out days and the current ~400-burger planning yield, that is roughly **2,900 DKK/week** of additional farmer payment. This is not a separate moral bonus outside the restaurant economics. It must be paid by the same whole-animal ledger as every other cost.

That is a gross-payment sensitivity only. Farmer net benefit also depends on:

- transport;
- delivery cadence;
- grade/quality requirements;
- rejected animals;
- payment terms;
- slaughter scheduling;
- additional handling;
- whether the direct relationship displaces a simpler alternative sale.

A higher gross payment that creates equal or greater new cost for the farmer is not success.

The processor has a similar hidden-cost boundary. Public custom-slaughter price lists are useful context, but a real 100BAD quote must include the actual specification: bespoke cutting, labeling, packaging, chilling, freezer handling, storage, scheduling, traceability and any small-batch surcharge.

³Danish Crown, *Kreaturnotering, uge 33 2026*. Used as a dated settlement benchmark, not a farmer contract. Danish Crown PDF

The whole-animal thesis then turns on **secondary absorption**. Under the balanced R1 routing scenario, the project models roughly **177 kg of non-premium edible output per reference animal**. At five sold-out days, 500 premium sales divided by the current ~400-burger planning yield requires about **1.25 animals/week**. That implies roughly **221 kg/week of secondary edible output** needing credible sales channels. Exact annual burden still depends on opening frequency and the observed reserve factor.

Three broad absorption channels are therefore modeled:

1. **Direct retail** - steaks, roasts and other higher-value cuts;
2. **B2B/frozen** - standard patties or other predictable secondary products;
3. **Fallback/low-realization** - channels that clear inventory but capture less value.

S60/S80/S100 are not profit margins. They are scenarios for how much of theoretical secondary value is actually realized after commercial friction.

The thesis fails if those products cannot be sold at sufficient realized value without excessive storage, labour, markdowns or working capital. It also fails if the higher farmer payment does not improve farmer net outcome.

THE WHOLE-ANIMAL BOTTLENECK


Five sold-out opening days imply about 1.25 animals/week and roughly 221 kg/week of secondary edible output.

K2 mass balance: in steady state, secondary outflow must keep pace with that generation over the allowed inventory horizon or inventory grows.

K2 value floor: the +15% farmer premium alone requires at least about 13.1 DKK/kg of *net* secondary contribution if secondary value is to fund that increment.

K3 farmer boundary: the gross premium is about 2,328 DKK/animal, so K3 fails if additional farmer-side obligations caused by the arrangement consume that amount or more.

These are quantitative floors, not evidence that the channels or farmer benefit exist.



Schematic, not to scale. Mixed EXTERNAL / SCENARIO / MODELED evidence.

Figure 4: Joint-product dependency. A higher farmer payment is economically plausible only if secondary outputs achieve enough realized value after processing and commercial friction.

5. Can 99 DKK Actually Work?

The reference customer proposition is **burger + tallow fries + house soda for 99 DKK including VAT**. This remains a stress-test price, not an observed optimum.

At the current VAT anchor, 99 DKK including VAT gives **79.20 DKK ex VAT** menu revenue per combo.⁴

⁴Skattestyrelsen, Danish VAT start guide; the general Danish VAT rate is 25%. Skattestyrelsen source

The previous reader draft then used a 28 DKK/combo food-cost placeholder, valued owner labour, and derived an approximate 70-sales/day K1 threshold. That was not a safe synthesis. The 28 DKK number was never decomposed clearly enough to tell whether it already contained the base beef cost, the +15% farmer premium, secondary-product credits or packaging.

The ~70/day claim is therefore **retired**.

One economic ledger

The model now uses one identity:

$$\text{premium menu revenue} + \text{net secondary contribution} - \text{whole-animal landed cost} - \text{non-beef food/packaging} - \text{owner labour} - \text{support} - \text{other variable cost} - \text{fixed overhead} = \text{economic surplus}$$

The +15% farmer payment lives inside **whole-animal landed cost**. It does not get paid a second time from a residual bucket. Secondary value can offset whole-animal cost only once.

The current whole-animal scenario uses:

- about **17.85k DKK/animal** farmer payment under the +15% case;
- public processing context of roughly **3.1-4.5k DKK/animal** ex VAT⁵;
- a midpoint landed-cost illustration of about **21.65k DKK/animal** before unresolved transport and other missing landed costs.

The non-beef food + packaging subtotal is not measured. Instead of hiding that uncertainty inside 28 DKK, the desk test now exposes **15 / 18 / 21 DKK per combo** as scenario inputs.

The coupled K1/K2 question

At 99 DKK, viability depends on how much **net secondary contribution after directly attributable secondary-channel costs** survives.

Using five opening days, modeled owner labour, the existing support/other-variable/fixed-overhead placeholders, midpoint processing context and the +15% farmer-payment case, the secondary contribution required for economic break-even is:

Table 2: Required net secondary contribution for weekly economic break-even. All cells are SCENARIO/D, not observed channel performance.

Sales/day	Non-beef 15	Non-beef 18	Non-beef 21
70	110.5	117.3	124.1
80	94.3	101.0	107.8
90	81.6	88.4	95.2
100	71.5	78.3	85.1

Values are DKK/kg net secondary contribution. Midpoint public processing context is used; unresolved transport/landed costs would raise the requirement.

At sold-out volume with the **18 DKK non-beef scenario**, the model requires about **78.3 DKK/kg net secondary contribution** to reach zero economic surplus.

That is the central reconciliation. The farmer-premium delta alone needs only about **13.1 DKK/kg** across the secondary pool, but the complete 99 DKK business needs substantially more because the same ledger must also cover the base animal, processing, human work and overhead.

The historical S60 scenario represents about **96.3 DKK/kg of realized ex-VAT secondary revenue after a 40% haircut to the benchmark basket**. It is not automatically net contribution because direct channel cost has not been observed. If future evidence showed that approximately that amount actually survives after

⁵Public processing-price context: Tistrup Slagtehus, prices effective 1 May 2026 (Tistrup source); Hallegaard Slagtehus, public cattle-processing price list (Hallegaard source). These are context examples, not directly comparable commercial quotes for 100BAD.

all directly attributable secondary-channel costs, the 18 DKK non-beef scenario would cross desk break-even around **83 sales/day**. That remains a conditional illustration, not a K1 pass.

Economic kill condition K1

K1 now fails whenever the combined menu + secondary ledger cannot cover:

- whole-animal landed cost including the farmer-payment scenario;
- non-beef food/packaging;
- owner labour at the chosen valuation;
- support;
- other variable operating cost;
- fixed overhead.

Current desk verdict: K1 OPEN / CONDITIONAL.

The model has not been killed at 99 DKK, but the paper no longer claims a single ~70/day break-even point. The actual threshold is a surface connecting utilization, secondary net contribution, non-beef cost and landed-animal cost.

CURRENT DESK RESULT

99 DKK remains mathematically possible in the current model, but only if the secondary stream creates enough *net* value. At 100 sales/day and an 18 DKK non-beef scenario, the current midpoint ledger requires about 78.3 DKK/kg net secondary contribution. K1 therefore cannot be separated from K2.




Figure 5: Unified K1/K2 stress test. Each line shows the net secondary contribution required for weekly economic break-even at 99 DKK under a different non-beef food/packaging scenario. The low dashed line is only the farmer-premium offset floor, not the business break-even threshold.

6. Can Two People Run It?

The reference uses **two operators** because robustness matters under bursty demand.

Dry Run 001 simulated 100 uneven orders across a four-hour **customer-facing service window**, including a short peak equivalent to 50 combos/hour. The reproducible assumptions and simulation output are included in the public implementation pack. The four hours are not the total workday. Prep, receiving, close, admin, cleaning and exceptions sit outside that window.

The two-person pipeline remained controlled under current touch-time assumptions. One-person scenarios

became fragile as integrated service time increased. This remains MODELED/C until Kitchen Pilot 001.

The finish/front model estimates about 53 seconds of active touch time per combo after plausible overlap. Ticket-machine or prepaid flow can reduce human payment/order-entry at the production station.

At 50 combos/hour, 150 g fries per combo implies **7.5 kg of finished fries per hour**. Reviewed commercial fryer capacities suggest raw equipment throughput is plausible, but sensory quality, recovery and tallow performance remain open.⁶

Cleaning remains counted as human work even when outsourced.

The current sold-out-day accounting is approximately:

- **10.4 combined owner/operator hours**
- **1.5 external cleaning hours**
- **11.9 total human hours**

That is the correct denominator for a human-scale claim.

The project now adds an explicit human-time kill condition:

Human-time kill condition K4

If measured total human work materially exceeds the bounded-work target without enough corresponding economics to justify the difference, the current human-scale architecture fails.

Future pilots must therefore measure not only throughput but:

- total operator hours;
- support hours;
- receiving/admin time;
- exception time;
- schedule fragmentation;
- holiday/sickness coverage pressure.

The cap creates a bounded production target. It does not magically create leisure.

Table 3: Modeled human-work boundary for a sold-out reference day.

Work boundary	Modeled time
Combined owner/operator work	10.4 h
External cleaning	1.5 h
Total human work	11.9 h

Evidence: MODELED/C. A physical pilot must replace these estimates.

⁶Capacity sanity check used manufacturer specifications from Frymaster RE80 fryers and Garland commercial griddles. These references establish only that relevant equipment classes exist; they are not procurement recommendations or observed 100BAD throughput. Full source notes are in the implementation pack evidence ledger.




Figure 6: Reference two-operator service pipeline. The 50-combo-per-hour peak remains a physical test target.

7. What Must the Local Network Provide?

At five sold-out opening days, the reference unit needs **500 sale buns per week** and **75 kg of finished fries per week**. These are useful because they make demand legible enough for upstream suppliers to quote.

The bun interface specifies functional properties rather than a brand. The potato interface can be whole-potato or local prep-node. Greens remain technology-neutral. The house-soda layer uses a small bounded syrup requirement.

The supplier rule remains:

If the supplier needs price X to be viable and 100BAD economics only tolerate much less than X , revise the restaurant architecture rather than assuming the supplier should absorb the gap.

This is not charity. It is system accounting.




Figure 7: Readable supplier interfaces. The cap converts vague local purchasing into bounded quantities, quality and cadence requirements.

8. Can Coordination Stay Thin?

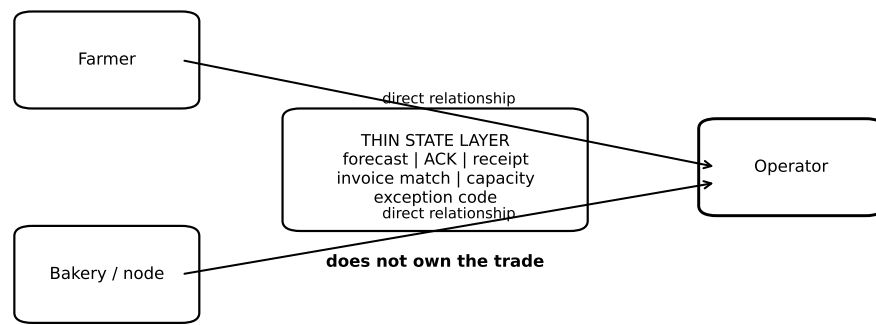
Local supply can recreate hidden labour through administration. Four small suppliers can mean repeated calls, invoice mismatches, late deliveries and constant exception handling.

The current design target is no more than **45 routine owner/operator coordination minutes per week** at five opening days. The event contract and coordination budget are included in the public implementation pack. The current event model lands close to that target, but the exact modeled minute count is low-confidence and belongs in the implementation appendix until observed.

Normal state should be automated or machine-readable: forecast, standing order, acknowledgement, receipt, inventory update and invoice match. Humans handle exceptions.

The coordination architecture fails its current labour target if routine owner time exceeds **90 minutes/week** absent genuine exceptions.

The shared layer may provide state without owning trade. The farmer and operator should be able to retain their direct relationship even if the software changes.



Design target ≤ 45 routine owner min/week; K6 if > 90 absent genuine exceptions.

Evidence: DESIGN/C + MODELED coordination logic.

Figure 8: Coordination without transaction ownership. Shared infrastructure carries operational state and exceptions while the commercial relationship remains direct.

9. Who Owns What?

The governance principle can be stated briefly:

No outside party should need to own the relationship between farmer and operator.

The restaurant belongs to the operator. The farm belongs to the farmer. Supplier nodes belong to their producers. Shared protocol/infrastructure can be stewarded separately.

The main paper does not claim that a Commons is necessary for restaurant viability. Governance is a secondary design layer **designed to reduce** future capture risk if the protocol scales; whether it actually does so is an open institutional hypothesis.

Concrete anti-capture requirements belong mainly in the implementation pack:

- data export;
- forkability;
- registry portability;
- trademark separation from transaction control;
- replaceable software maintainers.

Legal form remains OPEN and should not be treated as a solved part of the business model.

10. What Would Make Us Abandon It?

The workpack uses seven evidence classes: **OBSERVED, MODELED, EXTERNAL, SCENARIO, DESIGN, HYPOTHESIS, SPECULATIVE.**

Current stronger claims are structural or arithmetic:

- 100 sale tokens are a design rule;
- 50/hour with 150 g fries implies 7.5 kg/hour;
- five sold-out days imply 500 buns;
- suitable commercial equipment classes exist.

Current weak claims are commercial and empirical:

- roughly 400 premium patties per animal;
- 99 DKK viability;
- +15% farmer premium;

- customer demand;
- secondary absorption;
- actual coordination minutes;
- independent replication.

Kitchen Pilot 001 remains the next physical evidence upgrade: 20 sale tokens, forced short rush, timestamps, fryer recovery, sensory quality, reserve use, waste and operator stress.

Table 4: What would make the current model fail?

ID	Failure condition	What it kills
K1	Unified menu + secondary ledger cannot cover all economic costs	Reference economics
K2	Secondary outflow/value cannot meet mass-balance and ledger requirements	Whole-animal thesis
K3	Farmer incremental obligations consume the gross premium	Farmer-benefit thesis
K4	Total human work materially exceeds bounded-work target	Human-scale claim
K5	50/hour peak causes unacceptable quality/safety	Service architecture
K6	Routine coordination exceeds 90 min/week absent true exceptions	Distributed-node design

KILL CONDITIONS ARE NOT EQUALLY MATURE

K4 and K6 already have explicit design thresholds. Gate 25.3 adds quantitative floors to K2 and K3, but both remain empirically open. K2 still needs real channel sell-through, direct channel costs and an inventory horizon. K3 still needs a farmer-side cost quote. Numeric structure is not the same thing as evidence.

NOT YET PROVEN

No 100BAD restaurant is operating. No farmer contract, real utilization curve, physical 100-order day or replicated launch exists. M2 means the design is ready to be tested, not that the business has been demonstrated.

11. How to Try It

The accompanying **Public Implementation Pack RC1** exposes the mechanics rather than reserving them as private consulting material. It is a curated release surface, not the full historical research workpack.

It includes:

- evidence ledger;
- carcass and economics models;
- pilot protocol;
- supplier-node interfaces;
- coordination design;
- governance skeleton;
- replication checklist.

The reference independent launch sequence remains a **30-day design protocol**, not a guarantee.

The current training target remains HYPOTHESIS/D. The independence test remains whether a second competent operator can operate with minimal founder clarification and whether repeated questions become documentation fixes.

The invitation is:

If you try this, publish what breaks.

No explicit reuse license is assigned in RC1. The project intends an open implementation release, but the legal license must be chosen before the final public upload.

The design gains evidence whether the result is success or failure.

12. Would This Work Outside Denmark?

This case is calibrated to Denmark. The architecture may be portable; the current numbers are not.

A cross-country test would need to localize at least:

- menu price and consumption tax;
- labour valuation and employment structure;
- rent and fixed operating costs;
- cattle settlement and processing economics;
- food-safety and business-registration requirements;
- secondary-product demand and realized channel value.

The cap, sale-token architecture, whole-animal ledger, human-work accounting and kill-condition method can travel as hypotheses. **The 99 DKK result cannot.** External validity therefore remains a separate research question rather than a claim of Case 001.

A later comparative paper can test whether 100BAD is primarily a Danish configuration or a portable productive-unit architecture by holding the design logic constant and replacing local economic inputs.

PORTABILITY STATUS

OPEN. Case 001 does not claim that the Danish reference model works unchanged in Sweden, Germany, Poland, the United States or any other market.

Conclusion

100 Burger a Day is not mainly an argument about making burgers scarce. It is an attempt to choose a different accounting boundary.

The model counts:

- the animal;
- owner time;
- support labour;
- secondary outputs;
- supplier economics;
- utilization;
- coordination;
- governance risk.

That wider boundary makes the concept harder to flatter with selective arithmetic.

The model may still fail. The 99 DKK target may be too low. Secondary products may not clear. Farmer net benefit may disappear after added obligations. Total human hours may be too high. Real customers may not care.

Those are not embarrassing caveats. They are the next experiments.

**If the model is wrong, the next operator should be able to show us where.
If it is right, they should not need our permission to build it.**

Appendix A. Method and Evidence Grammar

100BAD uses a simple design-research loop:

1. choose explicit constraints;
2. model the productive unit;
3. expose hidden work and interfaces;
4. attack the model adversarially;
5. define kill conditions;
6. run the smallest physical test that can upgrade confidence;
7. revise or abandon the architecture.

The two headline constraints - 100 premium sale tokens and 99 DKK for the complete combo - are therefore not findings. They are stress tests. They force the model to reveal whether its economics depend on endless throughput, hidden labour, low supplier prices or unrealistic product routing.

Quantitative claims are tagged as OBSERVED, MODELED, EXTERNAL, SCENARIO, DESIGN, HYPOTHESIS or SPECULATIVE. A design choice can be high-confidence because it is deliberately chosen. A carcass yield can remain low-confidence even if the spreadsheet is internally consistent. Negative results are part of the method: if a pilot breaks the 50/hour service architecture or real quotes break the 99 DKK case, the result is recorded as evidence rather than repaired rhetorically.

References

External sources are cited where their claims appear. Internal figures, thresholds and design choices are reproducible from the accompanying Public Implementation Pack RC1.

1. Eurostat. “27.5% of self-employed people worked 49 hours or more per week in 2024.” 21 August 2025. Eurostat source
2. Mississippi State University Extension. “How Much Meat to Expect from a Beef Animal: Farm-Direct Beef.” MSU Extension source
3. Danish Crown. *Kreaturnotering, uge 33 2026*. Danish Crown PDF
4. Tistrup Slagtehus. Public processing prices, effective 1 May 2026. Tistrup source
5. Hallegaard Slagtehus. Public cattle-processing price list. Hallegaard source
6. Skattestyrelsen. VAT start guide. Skattestyrelsen source
7. Fødevarestyrelsen. Foodtruck/madbod rules and registration guidance. Fødevarestyrelsen source
8. Fødevarestyrelsen. Autorisationsvejledningen §15 on mobile/temporary retail food premises. Fødevarestyrelsen guidance
9. Frymaster. RE80 and Energy Star commercial fryer specifications. Manufacturer capacity context; source notes in the implementation pack.
10. Garland. Master Series / INSTINCT commercial griddle specifications. Manufacturer capacity context; source notes in the implementation pack.

SOURCE BOUNDARY

Public processor price lists are context, not 100BAD commercial quotes. Danish Crown is a settlement benchmark, not a farmer contract. U.S. carcass data is an aggregate technical sanity check, not an observed Danish carcass breakdown. Manufacturer specifications establish equipment-class plausibility, not observed service quality or throughput.