Report for IBIC funded project Phase 1 titled:

Automated detection and scoring of lameness and digital dermatitis lesions using artificial intelligence enhanced computer vision models upon harvesting of beef cattle

Submitted to: Iowa Beef Industry Council
Submitted by: Doerte Doepfer, DVM, MSc, PhD, professor

Department of Medical Sciences, School of Veterinary Medicine, 2015 Linden Drive, Madison, WI, 53706

Date: July 15th, 2026

Duration of project: October 2024 – May 2026 including 2 no cost extensions

NON-TECHNICAL SUMMARY

Lameness and its most common infectious cause, digital dermatitis (DD, also called hairy heel warts), are among the most costly and welfare-relevant health problems in U.S. dairy and beef cattle. Despite this burden, no automated system existed to measure how many cattle arrive at harvest lame or affected by DD, or to trace that information back to the producers of origin. This project built and deployed the first automated, camera-based artificial intelligence (AI) system to detect and score lameness and DD in beef cattle upon arrival at Upper Iowa Beef (UIB), a commercial harvesting plant in Lime Springs, IA.

Over nearly two years (October 2024 – May 2026), the project team collected more than a year of video footage across all seasons and lighting conditions and solved a dozen real-world engineering and biological problems — from lighting problems and camera-lens smudging to false cattle counts and curved walking paths — that stood between raw video and reliable AI-enhanced detections. The result is two fully operational AI pipelines, now running continuously on a secured onsite computer at UIB, that automatically detect lameness and DD and deliver a structured prevalence report for every truckload of cattle every 24 hours. The team also confirmed that a future AI chatbot could run on the same class of hardware, opening a path toward automated, interactive advice for cattle producers.

The single biggest limitation uncovered during Phase 1 is that cattle currently arrive at the harvesting plant without a readable individual identification, so detections cannot yet be linked to the specific animal. Solving this — together with connecting the system to UIB's producer-reporting software, quantifying the economic cost of lameness and DD, and extending monitoring back to the feed yards — is the focus of the proposed future

Phase 2 project, “Automated Lameness and Digital Dermatitis Surveillance at Harvest — Scaling to Identified Cattle, Feed Yards, and AI-Driven Producer Reporting.”

TECHNICAL REPORT

Impacts:

Phase 1 delivered the first scalable, automated surveillance infrastructure for lameness and digital dermatitis (DD) at the point of harvest in U.S. beef cattle (as of May 2026). By demonstrating that scalable computer vision (CV) and machine learning (ML) models can run continuously and unattended on edge hardware inside a working harvesting plant, the project moved lameness and DD monitoring from labor-intensive, infrequent manual scoring to continuous, load-level, 24-hour automated reporting. This establishes the technical foundation needed to eventually connect harvest-level production and health outcomes back to individual producers and feed yards of origin — a prerequisite for evidence-based, economically justified prevention programs across the Iowa beef supply chain.

List of quantitative impacts:

- Approximately 500 beef cattle assessed during preliminary data collection (November 2023) at the UIB unloading ramp; approximately 50% showed moderate-to-severe lameness and approximately 75% showed signs of active DD lesions.
- 460 feedlots within an approximately 150-mile radius of UIB supplied cattle between 2023 and 2026, providing load-level diversity in management systems and cattle types for model development.
- 2 computer-vision/machine-learning pipelines (lameness and DD) fully deployed and operational at UIB as of May 2026.
- Structured lameness and DD prevalence CSV reports generated automatically for every cattle load, every 24 hours.
- 12 major technical and biological challenges identified and resolved during Phase 1 (see Results).
- 1 edge-computing device (NVIDIA Jetson AGX Orin) deployed on site including systemd watchdog auto-restart in case of temporarypower loss, a prototype AGX Orin device has been hardened for cybersecurity with SSH, UFW firewall, AppArmor, and systemd watchdog auto-restart.

- LLM inference confirmed operational on NVIDIA Spark (DGX GB10) hardware, establishing feasibility for a future producer-facing AI advisory agent during Phase 2 in the future.
- Infographics and a prototype producer-facing reporting dashboard developed for extension and stakeholder communication (Figure 1).
- More than 12 months of continuous, multi-season video data collection completed to support model training.
- Data analysis of DD prevalence combined with cattle load harvesting data collected upon arrival at the harvesting plant.

Introduction

Identifying cohorts of beef cattle at high risk for health problems remains one of the most pressing challenges in U.S. beef production. Lameness and its primary infectious cause, digital dermatitis (DD, syn. hairy heel warts), and foot rot rank among the most prevalent and economically consequential conditions in feedlot and beef cattle systems (1–3). Morbidity from lameness, together with related production diseases such as bovine respiratory disease complex (BRDC) or gastrointestinal (GI) infections, results in severe production losses and animal welfare problems despite substantial research investment (1,2). Artificial intelligence (AI) and machine learning (ML) approaches have shown promising results for predicting high-risk cattle cohorts for disease outcomes, and AI-enhanced systems are increasingly implemented across animal agriculture (4–7). Computer vision (CV), a branch of AI that enables real-time analysis of video and image footage, is now applied to monitor cattle health and disease at scale (7–13).

The harvesting plant presents a uniquely accessible surveillance opportunity, because essentially all cattle pass through some form of unloading 'ramp,' which should make it possible to collect systematic, population-level data on lameness and DD that is, in principle, traceable back to the farm of origin. This opportunity had been systematically underutilized because no automated, scalable surveillance system existed prior to this project. Phase 1 (October 2024 – May 2026) was designed to close that gap by building and deploying deployable CV pipelines for lameness and DD detection at Upper Iowa Beef (UIB, Lime Springs, IA) — the first step toward a future system that reports lameness and DD prevalences and that identifies cattle at the farm of origin and links harvest-level detections back to individual producers.

Materials and Methods

Data collection

Preliminary video data were collected in November 2023, when approximately 500 beef cattle were observed during unloading at UIB using GoPro Hero 5 Black cameras during a high-risk period for DD in the Midwest. This footage confirmed that camera placement at the unloading ramp yields image quality sufficient to train accurate CV models and established the baseline prevalence used to justify the project (approximately 50% moderate-to-severe lameness, approximately 75% signs of active DD). UIB records further showed that 460 feedlots within an approximately 150-mile radius supplied cattle in 2023, providing substantial load-level variation across farms, management systems, and cattle breeds visualized by coat colors.

Building on these preliminary data, systematic video data collection for model development ran for more than 12 months to capture all seasons and lighting conditions encountered at the unloading ramp, and to resolve systematic misclassification of non-cattle objects (birds, tablets, paddles, people, boots) entering the camera field of view. Reliable, continuous video capture also required on-site infrastructure changes: LED lighting and a curtain were installed at the unloading ramp to stabilize lighting and shadows, and a self-cleaning ToughEye camera (ExcelSense, Canada, 15) was installed at floor level, together with a lens-cleaning routine needed to keep the camera functional in the particulate-rich slaughter-plant environment.

Approach to AI-enhanced estimation of DD and lameness prevalence

Two CV/ML pipelines were developed and deployed on a secured NVIDIA Jetson AGX Orin edge device (NVIDIA, Santa Clara, CA, 14) at UIB. The lameness pipeline links two CV models — cattle detection and cattle movement analysis — with a Random Forest ML classification model that scores lameness from movement features. The DD pipeline combines cattle detection with a DD detection model (16). Both pipelines use a multi-object tracker with 'virtual trackers' that follow individual cattle through the field of view of all cameras; this tracker was corrected during Phase 1 to eliminate systematic overestimation of cattle numbers ('ghost tracks'), and the underlying models were adapted to prevent gait-scoring artifacts caused by curved, multi-animal walking paths at the ramp (Figure 1c,d).

All model pipelines were reformatted to run on the Jetson AGX Orin (JetPack 6.2, TensorRT FP16), which required custom compilation of OpenCV 4.10.0 from source and resolution of several software incompatibilities. The deployed edge device needs to be hardened for cybersecurity (SSH, UFW firewall, AppArmor profiles, systemd watchdog auto-restart) to support unattended, continuous operation inside the plant. A completely hardened prototype of the system has been formatted, and test-runs have been performed before the end of Phase 1 of the project. In parallel, the team tested large language model (LLM) inference and a prototype chatbot on NVIDIA Spark (DGX GB10, aarch64) hardware, confirming the technical feasibility of a future AI advisory agent for producers during the future Phase 2 of this project.

Data analysis

Each pipeline outputs structured CSV files summarizing lameness and DD detections per animal and per cattle load, delivered automatically every 24 hours. Load-level DD and lameness prevalence are calculated directly from these structured outputs. Integration of these CSV outputs with UIB's partner reporting system, MeritTrac (Karnataka, India), and its SQL database was emphasized and technically prepared during Phase 1, laying the groundwork for automated producer-facing reporting, for example, using a dashboard (Figure 1e). To support producer and stakeholder communication of these results, the team also developed a series of infographics (Figure 1a,b) that translate the project's methods and data-collection context into accessible visual summaries for use during workshops and one-on-one producer conversations, so-called focus groups.

Harvesting Data Collection and analysis

Data were collected from 455 loads from 76 operations, from August 2024 until September 2025. Feet were scored on 269 of those loads (10,047 head) and combined with data after harvesting. Details of the anonymized data analysis are in supplements 1a (including the statistics) and 2 (written for producers without statistical jargon). Supplement 3 showcases one producer over time and is available only with permission of the producer.

Results

Phase 1 resulted in two fully deployed, operational AI pipelines for lameness and DD detection running on the Jetson AGX Orin edge device at UIB as of May 2026, generating structured CSV lameness and DD prevalence outputs per load, per day. Reaching this point required identifying and resolving 12 major technical and biological challenges, summarized below because they directly informed the objectives proposed for Phase 2:

1. Data collection required more than 12 months during Phase 1 to capture all seasons and lighting conditions, and to resolve systematic misclassification of non-cattle objects — birds, tablets, paddles, people, and boots — in the camera field of view, delaying full pipeline deployment.

2. LED lighting of the unloading ramp and installation of a curtain on the side of the unloading ramp were necessary for consistent light and shadows; repeated hardware and software failures caused pipeline downtime requiring repeated repairs.
3. A lens-cleaning routine for the ToughEye camera, installed at floor level close to walking cattle, was identified as a prerequisite for uninterrupted camera operation in the particulate-rich slaughter-plant environment.
4. Curved versus straight walking paths of multiple cattle arriving at once required model adaptation to prevent gait-scoring artifacts in the lameness pipeline (Figure 1c,d).
5. Systematic overestimation of cattle numbers (‘ghost tracks’) was corrected in the multi-object tracker of both pipelines using ‘virtual trackers’ that follow cattle through the field of view of all cameras.
6. All model pipelines were reformatted for the NVIDIA Jetson AGX Orin (JetPack 6.2, TensorRT FP16), requiring custom compilation of OpenCV 4.10.0 from source and resolution of software incompatibilities.
7. Experimentation with LLM inference and a prototype chatbot on NVIDIA Spark (DGX GB10, aarch64) confirmed the hardware and software feasibility required for a future AI advisory agent that is part of future Phase 2.
8. Individual cattle identification was found to be the main remaining challenge: without cattle IDs, pipeline outputs cannot be linked to individual animals or close-out groups from the farms of origin and their production data.
9. Without cattle identification, close-out based production losses attributable to lameness and DD cannot yet be quantified per animal or load, preventing evidence-based cost-benefit analysis.
10. Integration of pipeline CSV outputs with UIB’s partner, the MeritTrac API and SQL database, was emphasized and prepared during Phase 1 (Figure 1e); full integration and end-user reporting were identified as essential for producer adoption and is in the final process of implementation.
11. Edge device and pipeline security were implemented (SSH, UFW firewall, AppArmor, systemd watchdog) on a prototype AGX Orin device, establishing a security baseline for confident expansion to additional sites.
12. Infographics for workshops and extension communications were designed to convey complex information at a glance (Figure 1a,b), supporting producer focus-group conversations.

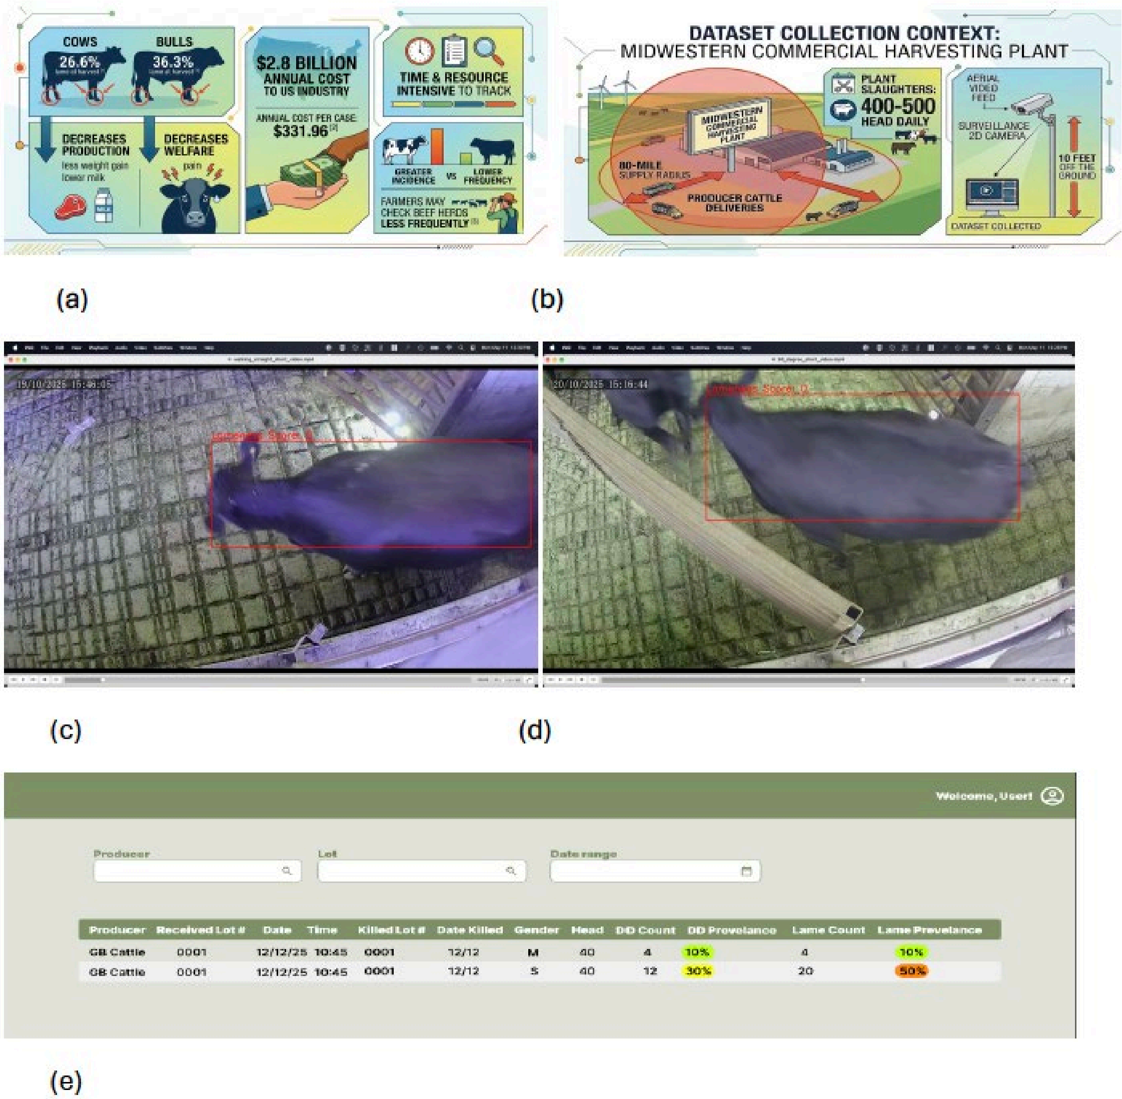

Figure 1 (a–e). (a,b) Infographics developed during Phase 1 to convey information about the beef cattle industry and UIB's producer clientele during focus-group and workshop conversations (infographics generated using ChatGPT.ai, used interactively as a design aid); (c,d) cattle walking in straight and curved paths at the unloading ramp, illustrating the gait-scoring challenge addressed in the lameness pipeline; (e) draft dashboard developed during Phase 1 for future interaction with end users during automated quarterly reporting.

Results from DD prevalence and harvesting data (see details in supplement 1a with statistics and anonymized text, figures and tables, supplement 2 written for anonymized producers, supplement 3 is the report for one producer serving as show case):

This report covers 455 loads from 76 operations, August 2024 to September 2025. Feet were scored on 269 of those loads (10,047 head).

1. **Data recovered.** 76 producers after name cleaning (from 93 raw factor levels); the prevalence model now uses **269 lots / 48 producers** instead of the 49 lots the old Origin variable allowed.
2. **DD burden.** 982 DD cases in 10,047 animals, pooled prevalence **9.8%**; 16% of lots have no DD at all.
3. **Producer matters.** Producer variance 0.47 on the logit scale (ICC 0.12), i.e. about 12% of the variation in DD risk sits between producers - the random id term is doing real work.
4. **Season.** Model-based prevalence: Spring 6.3%, Summer 4.8%, Fall 28.8%, Winter 4.6% (5 Feb lots only). Fall is far above every other season (all Tukey $p < 0.001$).
5. **Data corrected.** 28 mistyped values in 27 lots repaired (each explained by one keystroke slip), 6 values deleted as unreconstructable, 1 lot left for manual review; that returns 262 lots to the weight/yield models.
6. **Overdispersion.** Pearson chi2/df = 2.91; the beta-binomial GLMM fits far better (AIC 1266 vs 1501), so its wider intervals are the honest ones.
7. **Harvest outcomes.** After adjustment for Season, Gender and producer, and False Detection Rate (FDR due to multiple statistical testing) correction across the 15 outcomes, the surviving association is yield_pct_clean (0.174, q = 0.037). This means: with +10 percent DD prevalence, the yield_pct_clean increases by 0.174 units according to the data collected. This might be due to the shrinking of the carcass (in the denominator of the yield_pct-clean) not the gain of bodyweight.

Discussion and Outlook

Phase 1 demonstrated that automated, AI-enhanced surveillance of lameness and DD is technically achievable at commercial line speed and in the particulate-rich, variable-lighting environment of a working harvesting plant — an outcome not previously achieved at this scale. The preliminary and Phase 1 data confirm that lameness and DD remain highly prevalent in cattle arriving at harvest, underscoring the practical value of continuous, automated monitoring over infrequent manual assessment.

The central limitation identified during Phase 1 for connecting harvesting data to production and health data from farms of origin is the absence of individual cattle identification at the unloading ramp: without RFID or another reliable visual ID, load-level detections cannot be connected to individual animals, close-out groups per feed yard of origin, and no cost estimation, or predictive modeling can be pursued with certainty. Related limitations include the current disconnect between pipeline outputs and UIB's MeritTrac reporting system (technically prepared and in final stages of implementation), and the fact that the lameness and DD models were trained specifically on UIB footage and may require adaptation before they generalize to new camera angles, lighting conditions, or cattle populations at other locations. Generalizability of the detections has not be tested in the field during Phase 1 of the project.

These findings directly motivate the proposed Phase 2 project, “Automated Lameness and Digital Dermatitis Surveillance at Harvest — Scaling to Identified Cattle, Feed Yards, and AI-Driven Producer Reporting” (proposed period: August 2026 – July 2027). Phase 2 is designed to close the gaps identified in Phase 1 through eight objectives: (1) RFID-based or digital reading of real IDs for cattle identification at the unloading ramp; (2) integration of pipeline outputs into MeritTrac for automated quarterly producer reports; (3) predictive modeling linking lameness and DD detections to production and harvest outcomes, including partial-budget economic analysis (17,18); (4) semi-supervised feature learning to generalize the pipelines beyond UIB; (5) an AI advisory agent on NVIDIA Spark hardware, built on retrieval-augmented generation (22) and the LLM feasibility work confirmed in Phase 1 (19–21); (6) extension of cybersecurity features across edge devices; (7) deployment of the camera-pipeline system at producer and feed yard sites, upstream of harvest; and (8) continued focus groups, workshops, and dissemination with producers, veterinarians, and Iowa State University Extension personnel.

The application for IBIC funding for the 2026 cycle has been rejected, and therefore, Phase 2 has to seek different funding sources to guarantee continuity.

Anticipated challenges for Phase 2 include variable RFID tagging compliance across client feed yards, dependence on MeritTrac API access from the database vendor, the need for repeated model tuning as the system generalizes to new beef cattle operations, and the hallucination risk inherent to LLM-based advisory tools, which will be mitigated using a retrieval-augmented generation architecture grounded in expert-reviewed sources.

Taken together, Phase 1 established the operational core of an automated lameness and DD surveillance system; Phase 2 is designed to transform that core into a decision-support system that connects harvest-level detections to individual producers, quantifies the economic burden of lameness and DD, and extends monitoring upstream to the farms and feed yards where prevention decisions are made.

The data analysis resulted in a list of recommendations (for details consult supplement 2):

1. **The point of interest for change is the yard, not the season or the sex.** Roughly 12% of the variation of DD prevalence sits between operations. Some operations in this study run near 2% year-round; that is achievable and one can learn from those operations.
2. **Judge yourself on a running average, not on one load.** A 38-head load at 5% is one or two animals, and that number bounces for reasons that have nothing to do with your management. Ten or more loads give a figure worth acting on.
3. **Treat any load over about 20% as an event.** Write down the pen, the source group, how long they had been on feed and what the pen conditions had been. In this study, a small number of loads like these carried a large share of the total burden, and they are the cases with something to change at the yard of origin.
4. **Do not chase dressing percentage.** If it rises alongside more DD, that is shrink, not gain.
5. **Record three more fields at delivery** — placement date or days on feed, arrival weight, and source group. They cost nothing to capture and they would settle the biggest open question in this analysis.

Deliverables

- Two fully deployed, operational computer-vision/machine-learning pipelines (lameness and DD) running continuously on a secured edge device at UIB.
- A secured NVIDIA Jetson AGX Orin edge-computing installation (SSH hardening, UFW firewall, AppArmor profiles, systemd watchdog auto-restart).
- Structured, automatically generated CSV lameness and DD prevalence reports per cattle load, delivered every 24 hours.
- A documented record of the 12 major technical and biological challenges encountered and solved during Phase 1, forming the technical basis for Phase 2 planning. Solving these 12 challenges was not trivial, and the argument that Phase 1 had delayed delivery of outcomes ignores that additional goals have been achieved among which the cybersecurity and AI agent prototypes.
- Confirmed technical feasibility of LLM inference and a prototype chatbot on NVIDIA Spark hardware, supporting the design of a future producer-facing AI advisory agent.
- A set of infographics for producer and stakeholder communication, and a draft dashboard prototype for future automated producer reporting (Figure 1).
- Technical groundwork and final progress for integrating pipeline outputs with UIB’s MeritTrac reporting system.
- A Phase 2 research proposal, “Automated Lameness and Digital Dermatitis Surveillance at Harvest — Scaling to Identified Cattle, Feed Yards, and AI-Driven Producer Reporting,” submitted for continued funding.

In addition, there are **3 supplements with details about the statistical analysis** of the DD prevalence data combined with harvesting data.

Acknowledgements

Many thanks to the Iowa Beef Industry Council for funding this project and to the many producers whose harvesting data were collected. In addition a huge thank you to the team at Upper Iowa Beef Inc, in Lime Springs, IA for all patience and collaboration during the project.

The author would be happy to present the details of this report and outcomes to anyone interested in the matter.

References

1. Peel DS. The Effect of Market Forces on Bovine Respiratory Disease. Vet Clin North Am Food Anim Pract. 2020;36(2):497-508.

2. Wilson BK, Step DL, Maxwell CL, et al. Effect of bovine respiratory disease during the receiving period on steer finishing performance and carcass characteristics. Prof Anim Sci. 2017;33(1):24-36.

3. Kulow M, Merkatoris P, Anklam KS, Dopfer D, et al. Prevalence of digital dermatitis and effects on performance in beef feedlot cattle. J Anim Sci. 2017;95(8):3435-44.

4. Wisnieski L, Amrine DE, Renter DG. Predictive Models for Weekly Cattle Mortality after Arrival at a Feeding Location. Pathogens. 2022;11(4):473.

5. Wisnieski L, Amrine DE, Renter DG. Predictive modeling of bovine respiratory disease outcomes in feedlot cattle: A narrative review. Livest Sci. 2021;251:104666.

6. Beauchemin KA. Current perspectives on eating and rumination activity in dairy cows. J Dairy Sci. 2018;101(6):4762-84.

7. Borges Oliveira DA, et al. A review of deep learning algorithms for computer vision systems in livestock. Livest Sci. 2021;253:104700.

8. Eli-Chukwu NC. Applications of Artificial Intelligence in Agriculture: A Review. Eng Technol Appl Sci Res. 2019;9(4):4377-83.

9. Cernek P, Bollig N, Anklam K, Dopfer D. Detecting digital dermatitis with computer vision. J Dairy Sci. 2020;103(10):9110-5.

10. Higaki S, et al. Leveraging computer vision-based pose estimation in dairy cows for objective mobility analysis. Comput Electron Agric. 2024;217:108573.

11. Kang X, Zhang XD, Liu G. Accurate detection of lameness in dairy cattle with computer vision. J Dairy Sci. 2020;103(11):10628-38.

12. Kang X, Zhang XD, Liu G. A Review: Computer Vision-Based Lameness Detection for Dairy Cows. Sensors (Basel). 2021;21(3):753.

13. CattleEye. Autonomous Livestock Monitoring [Internet]. Available from: https://cattleeye.com/

14. NVIDIA. Jetson AGX Orin for Embedded & Edge Systems [Internet]. Available from: https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-orin/

15. Excelsense Technologies. ToughEye 1700 Self Cleaning Camera [Internet]. Available from: https://www.excelsensetechnologies.com/

16. Anklam K, Cernek P, Yoon I, Dopfer D, et al. Saccharomyces cerevisiae fermentation product for prevention and control of digital dermatitis. Appl Anim Sci. 2022;38(2):98-109.

17. Ahmed H, et al. Assessing Animal Welfare and Farm Profitability with Stochastic Partial Budgeting. Animals. 2021;11(2):382.

18. Ahmed H, et al. Economic consequences of improved animal welfare in cattle fattening using stochastic partial budgeting. Livest Sci. 2020;232:103920.

19. Touvron H, et al. LLaMA: Open and Efficient Foundation Language Models. arXiv:2302.13971; 2023.

20. Google. Gemma-7b. Hugging Face [Internet]. 2024. Available from: https://huggingface.co/google/gemma-7b

21. OpenAI. ChatGPT [Internet]. Available from: https://chat.openai.com/

22. Lewis P, Perez E, Piktus A, Petroni F, Karpukhin V, Goyal N, et al. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. Advances in Neural Information Processing Systems (NeurIPS). 2020;33:9459–9474. arXiv:2005.11401

23. HubSpot. Craft Your Own Python AI ChatBot. 2024.

24. Raschka, S., Building LLMs from the Ground Up: A 3-hour Coding Workshop. https://www.youtube.com/watch?v=quh7z1q7-uc

What this is
How much DD is out there
The biggest differences are between operations, not between seasons
Season
Heifers and steers
Does DD show up on the kill sheet?
The dressing-percentage puzzle
What I would take from this
How to read the numbers here

Digital dermatitis in fed cattle: what 455 loads show

A summary for beef producers - all operations, identities withheld

Prepared by D. Döpfer

01 September 2026

What this is

Every load of fed cattle delivered to one plant over about a year was scored for digital dermatitis (DD, hairy heel wart) before slaughter, and the score was matched to that load's own kill sheet — weights, dressing percentage and grade mix. This report covers **455 loads** from **76 operations**, August 2024 to September 2025. Feet were scored on **269 of those loads** (10,047 head).

No operation is named. Each appears only as a code - P01, P02 and so on - assigned alphabetically and meaningless outside this report. If you want to know which code is yours, ask the veterinarian or the plant; the key is held separately and is not part of this document.

The purpose is simple: to tell you how common DD is in cattle like yours, how much it varies from one operation to the next, and whether it shows up anywhere on the kill sheet.

How much DD is out there

DD across all scored loads

Measure	Value
Loads scored	269
Head scored	10,047

Measure	Value
Head with a DD lesion	982
Overall prevalence	9.8% of head
Typical (median) load	5.9% of the load
Loads with no DD at all	16% of loads
Worst load recorded	66% of the load

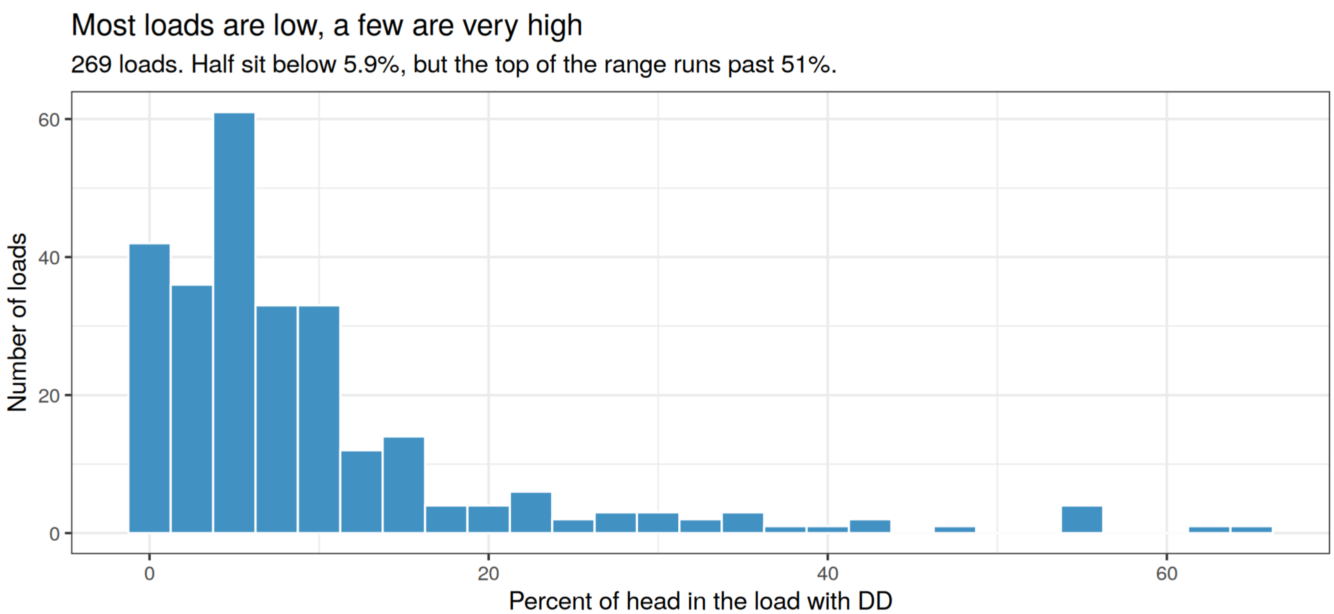

About one load in six arrives clean. Half of all loads sit below 6%. But the tail is long, and a handful of loads come in with a third to two-thirds of the animals affected — those loads are where most of the total burden sits.

The biggest differences are between operations, not between seasons

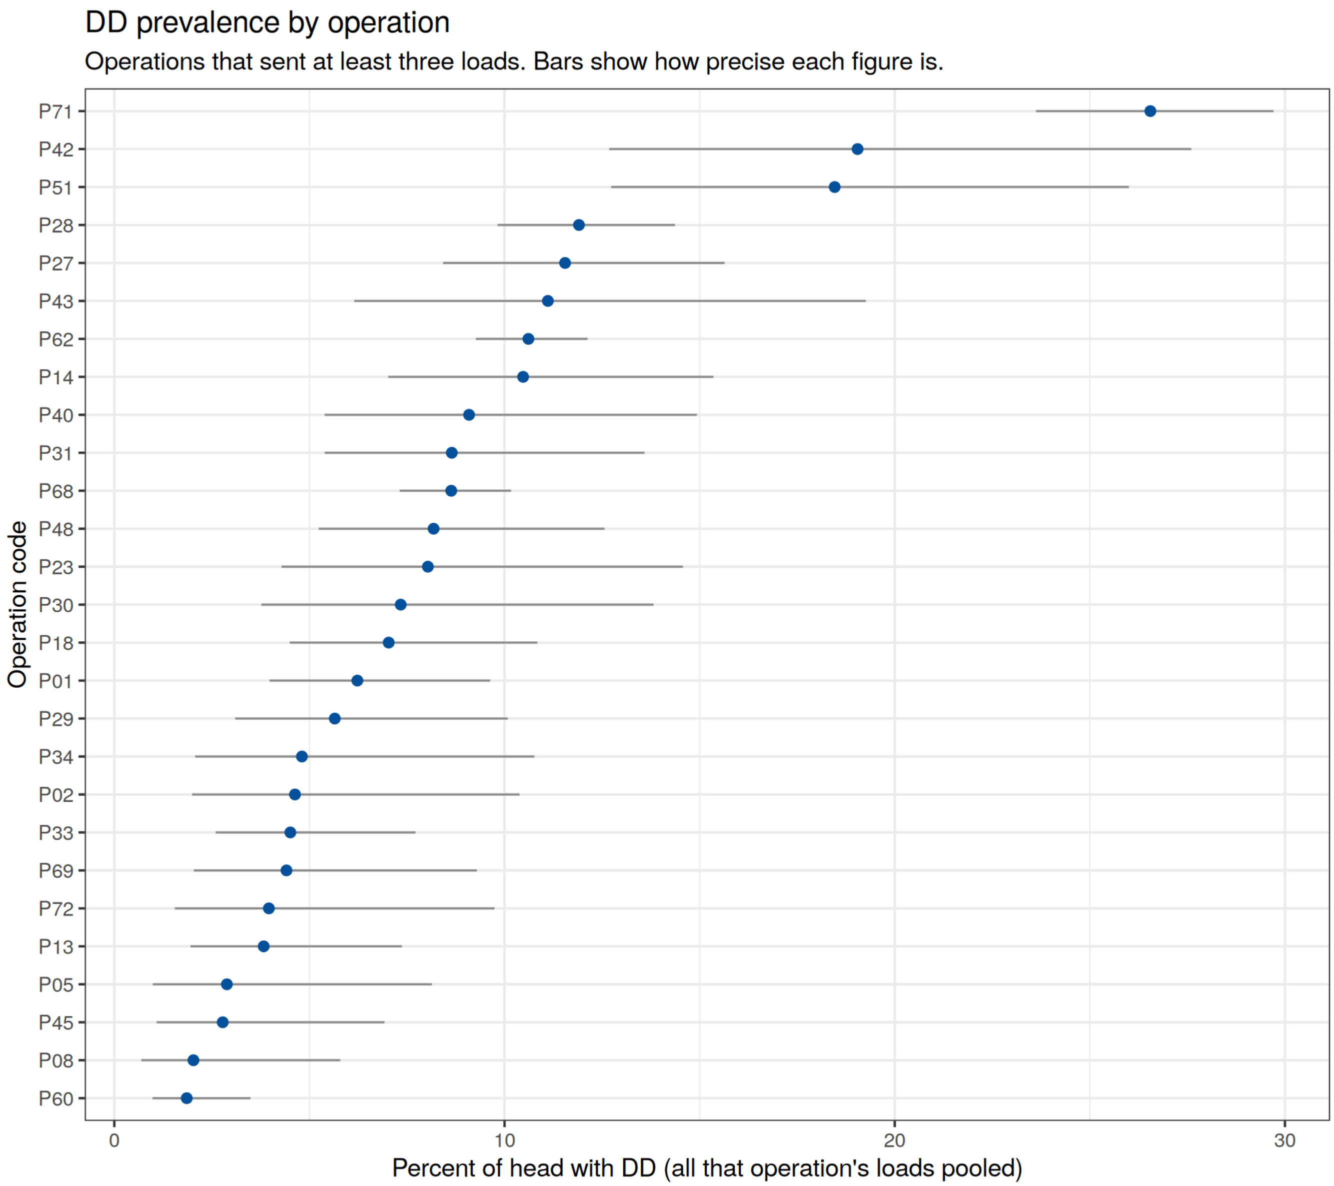

Spread across operations sending three or more loads

Measure	Value
Operations shown	27
Lowest	1.9% - P60
Median	7.3%
Highest	26.6% - P71
Share of the variation that sits between operations	about 12%

This is the most useful number in the report for a producer. Operations at the bottom of that chart run near **2%**; operations at the top run above **27%** — a 14-fold difference in cattle going to the same plant in the same year. Roughly **12%** of all the variation in DD risk is *between* operations rather than between individual loads.

Whatever is driving that — pen surface and drainage, bedding, stocking density, footbath use, source of the cattle, how long they stand on concrete — it travels with the operation, not with the season or the load. Which means it is largely under management control.

Season

DD prevalence by season, averaged across operations

Season	Loads scored	Predicted DD prevalence (%)	Range (95%)
Spring	193	6.3	5.0 - 7.9
Summer	61	4.8	3.6 - 6.3
Fall	10	28.8	20.6 - 38.7
Winter (Feb only)	5	4.6	2.5 - 8.3

Autumn stands out sharply — roughly 29% against 5-6% the rest of the year, and the difference is far too big to be chance. But read it with two cautions:

- it rests on only **10 autumn loads**, so a couple of bad consignments carry a lot of weight in that figure;
- **winter is nearly missing.** No cattle at all were delivered in November, December or January across the whole study, and "winter" here means 5 loads from a single February. The report cannot tell you what winter looks like.

So: autumn appears to be the risk period in these data, and that is worth a second year of records before anyone builds a program around it.

Heifers and steers

DD prevalence by sex, adjusted for season and operation

Cattle	Predicted DD prevalence (%)	Range (95%)
Heifers	5.1	3.8 - 6.7
Steers/bulls	6.4	5.1 - 8.1

A small difference, and not one to manage around.

Does DD show up on the kill sheet?

Fifteen things on the kill sheet were tested against how much DD a load carried — the three weights, dressing percentage, the five yield grades and the six quality-grade categories. Everything is adjusted for season, sex and which operation the cattle came from, and corrected for the fact that testing fifteen things at once will throw up a false positive on its own.

Effect of an extra 10 percentage points of DD in a load. Grade rows are multipliers (1.00 = no change).

Outcome	Loads	Change per +10% DD	Clear signal?

Outcome	Loads	Change per +10% DD	Clear signal?
Dressing percentage (points)	210	+0.17	yes
Share Select	181	x1.41	no - disappears once all 15 tests are allowed for
Share grading YG2	184	x0.88	no
Share grading YG1	184	x0.83	no
Share G1-37 upper two-thirds	181	x1.13	no
Share grading YG5+	184	x1.12	no
Share grading YG3	184	x1.03	no
Share CAB/Prime	181	x0.96	no
Share CAB upper two-thirds	181	x0.97	no
Live weight per head (lb)	208	-3.95	no
Share Choice	181	x1.04	no
Share grading YG4	184	x1.02	no
Carcass weight per head (lb)	207	+1.30	no
Shrunk weight per head (lb)	209	-1.68	no
Share G1-37 Prime	181	x1.03	no

The headline is a negative one, and it is good news as far as it goes: carcass weight, the grade mix and the marbling categories do not move with DD. A load with 20% DD hangs about the same carcass, in about the same grades, as a load with none. Only one measure moves clearly — dressing percentage — and it moves *up*, which needs explaining because it is not the bonus it looks like.

The dressing-percentage puzzle

The number that moved is the bottom of the fraction

Dressing percentage is carcass weight divided by live weight. It can rise because the carcass got heavier, or because the live weight got lighter. Splitting the two:

Where the dressing-percentage change comes from

Measure	Change per +10% DD	Solid enough to bank on?
Live weight at the plant	-4.1 lb/head	no - well inside the noise
Carcass weight on the rail	+0.8 lb/head	no - essentially zero
Dressing percentage	+0.20 points	yes - the clearest signal in the study

Put those back through the fraction and about **77% of the gain comes from live weight falling, not carcass weight rising.** No extra beef is being produced. The cattle are simply crossing the scale lighter and hanging the same carcass. Both weight figures are individually too uncertain to bank on, so treat that split as the direction of travel rather than a settled number — but the direction is clear, and it is not in your favour.

Why lighter alive but the same on the rail

The difference between an animal's live weight and its carcass is hide, head, feet, blood and — by far the most variable piece — **gut fill**, which runs roughly 6–12% of live weight and swings with how recently the animal ate and drank. Anything that empties cattle out lowers live weight, leaves the carcass alone, and pushes dressing percentage up on paper.

Three ways DD plausibly does that:

1. **Less time at the bunk and the tank.** Sore heels mean more lying, less walking, lower intake in the last day or two before shipping.
2. **Slower loading and more handling.** Lame cattle take longer up the alley, which means more time off feed and water, which means more shrink — mostly gut fill and body water.
3. **Longer standing at the plant.** Same story again at the other end of the haul.

Three explanations that were tested and do not hold

- **"It's the pencil shrink."** These records carry a shrunk weight, but it is a contract deduction (0%, 1% or 3% off live weight), not a second trip across the scale — across the whole study it takes only 13 distinct values in 262 loads. Adjusting for which deduction applied does not change the result at all, and working dressing percentage straight off live weight, where no deduction exists, makes it larger.
- **"Muddy hides."** Mud and tag add live weight and would push dressing percentage *down*. A dirty-pen explanation predicts the opposite of what we see.
- **"It's just chance."** Possible, but this is the strongest signal in the study and it survives the correction for testing fifteen outcomes at once.

The one that cannot be dismissed: time on feed

The sharpest objection is that cattle standing on feed longer get more exposure to DD **and** carry more finish, and finished cattle dress higher. Days on feed could drive both numbers with no connection between them at all.

No record in this study carries days on feed, placement date, arrival weight or age, so the question cannot be settled here. Probing it with the finish measures that *are* on the kill sheet:

- mean yield grade rises by 0.033 per +10% DD — the direction the time-on-feed story needs, but small and inside the noise;
- live weight does *not* rise, though that proves little, since cattle are marketed to a target endpoint and out-weights are held similar by design;
- adjusting dressing percentage for yield grade, marbling and live weight together moves it from 0.20 to 0.17 points — the measurable finish signal carries only about 12% of the result.

For time on feed to explain the whole thing, the part of it not visible in the grades would have to be several times more influential than the grades themselves. That is possible. It is not settled. Treat the dressing-percentage finding as a real pattern with one unresolved rival explanation.

What it means for your cheque

- **Selling live weight?** This runs against you — the estimate is about 4 lb less live weight per head for every 10 percentage points of DD. Small, not statistically firm, and pointed the wrong way.
- **Selling on a grid or rail weight?** Close to neutral. Carcass weight and grade mix did not move.
- **Either way, there is no premium hiding in a higher dressing percentage** that comes from cattle arriving empty.

And note what a kill sheet cannot see at all: treatment and labour, animals held back or railed, trim at the plant, and the welfare cost itself. The carcass record is a narrow window on what lameness actually costs an operation.

What I would take from this

1. **The lever is the yard, not the season or the sex.** Roughly 12% of the variation sits between operations. Some operations in this study run near 2% year-round; that is achievable.
2. **Judge yourself on a running average, not on one load.** A 38-head load at 5% is one or two animals, and that number bounces for reasons that have nothing to do with your management. Ten or more loads give a figure worth acting on.
3. **Treat any load over about 20% as an event.** Write down the pen, the source group, how long they had been on feed and what the pen conditions had been. In this study, a small number of loads like these carried a large share of the total burden, and they are the cases with something to teach.
4. **Do not chase dressing percentage.** If it rises alongside more DD, that is shrink, not gain.
5. **Record three more fields at delivery** — placement date or days on feed, arrival weight, and source group. They cost nothing to capture and they would settle the biggest open question in this analysis.

How to read the numbers here

- **DD prevalence** is the percentage of head in a load with a digital dermatitis lesion, scored at the plant before slaughter. Cattle were killed the same or the next day, so the foot score and the kill sheet describe the same animals on the same day.
- **The bars on the charts** show the range in which the true figure most likely sits. Wide bars mean few animals and an imprecise number.
- **“Adjusted for season, sex and operation”** means each comparison is made between similar loads — two loads from the same operation, in the same season, of the same sex — so that a producer who happens to run both high-DD and high-yielding cattle cannot create a false link.
- Everything here is an **association**, not proof of cause. Nothing in this study shows what would happen if DD were reduced — only what tends to accompany it.

Based on 455 delivered loads from 76 operations, 269 of them foot-scored. Weight records were checked for keying errors before analysis: 28 mistyped values in 27 loads were corrected, 6 values were removed as unusable, and 1 load was set aside for manual review. All operations appear as codes only.

1 Analysis strategy

2 Data-quality diagnosis and the `id` fix

3 Descriptive analysis

4 Bivariate analysis

5 GLMM: DD prevalence with random producer

6 Predicting harvest outcomes from DD prevalence

7 Summary of findings

8 Verification

UIB / DD: lot-level DD prevalence and its association with harvest outcomes

Doepfer

2026-09-01

1 Analysis strategy

Unit of analysis. One row = one *received lot* (n = 455). Lots are nested in **producers**; producer is the clustering (random) factor throughout. There is no repeated measurement *within* a lot, so a lot-level random effect is not identifiable separately from residual/overdispersion.

Primary outcome. DD prevalence in a lot = `DD_count / Head` (successes / trials), modelled as a binomial proportion — never as a raw percentage in a linear model.

Secondary (harvest) outcomes. Two families with different natural scales:

Family	Variables	Scale	Model
Weights / yield	`Total_Live_Weight` , `Shrunk_Weight` , `Dressed_Weight` , `Yield_pct`	**lot totals** → converted to **per head**	LMM (`lmer`)
Carcass counts	`YG1` – `YG5more` , `CAB_Prime` , `G137_Prime` , `Cab_Upper_2_3rd` , `G137_Upper_2_third` , `Choice` , `Select`	**counts of carcasses** out of the graded lot	binomial GLMM (`glmer`)

The weight variables are *lot sums* (mean live weight ≈ 56,000 lb per ~38 head). Modelling them untransformed would mostly model lot size, so every weight is divided by `Head` . The grade variables are *counts*, not percentages — they sum to the lot size — so they are modelled as proportions with the graded-lot denominator, not as Gaussian outcomes.

Analysis sequence

1. Clean → build a **producer-based `id`** (the fix described in §2).
2. Data-quality audit and **correction of the weight entry errors**, with a line-by-line record of every value changed (§2.3); missingness; why Winter is thin.
3. Complete descriptive analysis (lot level, producer level, by season / gender).
4. Bivariate analysis: each candidate predictor vs. DD prevalence; DD prevalence vs. each harvest outcome (unadjusted), with FDR control across the outcome family.
5. **GLMM for prevalence:** `cbind(DD_count, Head - DD_count) ~ Season (+ Gender) + (1 | id)` . DD counts turn out to be strongly overdispersed relative to the binomial (Pearson χ^2/df ≈ 2.9), so the **beta-binomial GLMM is the model used for**

inference and the plain binomial fit is kept only for comparison — reporting the binomial p-values alone would overstate precision. Reported: ICC, EMMs, Tukey contrasts, producer caterpillar plot.

6. **Harvest-outcome models:** `outcome ~ DD_prevalence + Season + Gender + (1 | id)`, one model per outcome from a single fitting function, results pooled into one table with FDR-adjusted p-values and effects expressed per **+10 percentage points** of DD prevalence.
7. Automated verification (assertions) at the end.

What changed relative to `DD6_UIB_predictions8_31_26.Rmd` — see §2.

Companion script. `DD7_build_clean_data.R` applies exactly the cleaning in §2 and saves `DD7_df_clean_8_31_26.RData` (objects `df`, `producer_lookup`, `codebook`). Use it when you want the cleaned data without knitting this report; this report still cleans from the original `DD5` file, so the two never drift apart.

2 Data-quality diagnosis and the `id` fix

2.1 What was wrong

Root causes found in DD6_UIB_predictions8_31_26.Rmd

Issue	Evidence	Consequence
id built with LETTERS[1:n]	Producer had 93 factor levels but only 26 letters exist	Origin was NA for 363 of 455 rows (80%)
Producer strings never normalised	trailing blanks, plus three spellings of one farm name and two of another	93 raw levels for 81 real producers
glmer dropped every row with NA id	complete cases for cbind(DD_count, Head) ~ Season + (1|Origin): n = 49	model was fitted on ~11% of the data
NA dates fell into the else branch	the old nested ifelse sent any row with a missing Date_Received to '4_Fall'	2 lots with no receive date were counted as Fall (39 vs the correct 37)
Winter 'missing'	Winter lots in raw data: 5 (all Feb-2025); with non-missing Origin: 0	level silently dropped from the design matrix

Why Winter was missing — two independent reasons.

1. *Mechanical:* all 5 Winter lots had `Origin = NA` (their producers fell past the 26th factor level), so `glmer` deleted them together with 358 other rows. A factor level with zero remaining observations is dropped from the model frame without a warning.
2. *Structural:* there genuinely is almost no winter in this dataset. Lots were received Aug-2024 → Sep-2025 but **no lot at all was received in November, December or January**; "Winter" consists only of five February-2025 lots, and all five have `Gender = NA`. So even with the id fixed, Winter cannot support a `Season × Gender` adjusted estimate. This file therefore reports Winter in the Season-only model and drops/collapses it in Gender-adjusted models (see §5.4), rather than pretending the level is estimable.

Lots received per calendar month: Nov-2024 - Jan-2025 are empty

ym	n
2024-08	2
2024-09	0
2024-10	10
2024-11	0
2024-12	0
2025-01	0
2025-02	5
2025-03	34

ym	n
2025-04	64
2025-05	105
2025-06	76
2025-07	76
2025-08	54
2025-09	27
NA	2

2.2 The fix: a producer-based `id`, and the weight corrections

All cleaning lives in **`DD7_data_prep.R`**, sourced here and by `DD7_build_clean_data.R`, so the report and the saved data set cannot drift apart. The producer index is built by **normalise** → **merge known spelling variants** → **index**, using `sprintf("P%02d", ...)`, which cannot run out of labels the way `LETTERS` did.

```
## Producer levels: 93 raw -> 76 after cleaning;  id missing: 0 rows
```

```
## Producer names are withheld in this version; producers appear only as id codes.
```

Producer -> id lookup (15 largest; full table in `producer_lookup`)

id	Producer_clean	n_lots
P62	P62	77
P68	P68	49
P28	P28	37
P71	P71	31
P01	P01	25
P60	P60	13
P27	P27	9
P45	P45	9
P66	P66	9
P18	P18	8
P36	P36	8
P40	P40	8
P42	P42	8
P31	P31	7
P33	P33	7

2.3 Correcting the weight entry errors

The three weight columns are lot totals typed by hand; `Yield_pct` comes from the packer and reproduces `Dressed/Shrunk` to about 0.01 points in the lots that are internally consistent. That makes the yield a **check digit**: where a lot fails, the recorded yield says which weight is wrong and what it should have been.

The rule applied is deliberately narrow — *repair the typing, never invent a measurement*:

1. A weight is replaced only by a value reachable from what was actually recorded through **one simple typing slip** — a decimal shift, one digit inserted, deleted, substituted, or two adjacent digits transposed.
2. That candidate must put the whole row back inside the plausibility envelope: per-head weight, shrink ratio `SW/LW`, and the recorded yield to within 0.5 points.
3. It must be the **only** such reading. If two different corrections would both work, nothing is changed and the lot is listed for human review.
4. A value that no typing slip explains is **set to `NA`** — deleted, not guessed. Rebuilding a dressed weight as `SW × Yield` would carry no independent information and would bias any model that uses those same quantities.

Because the repaired figure is the corrected *reading* rather than `SW × Yield`, each lot keeps its own small measurement error instead of being forced into perfect arithmetic agreement with the yield it is later compared against. The originals stay in `df` as `*_orig`, so every change is reversible.

Every corrected data point (n = 28 values in 27 lots)

Lot	Head	Variable	Recorded	Corrected	Slip identified
L041	15	Total_Live_Weight	193400.0	19340.0	LW: decimal shift x0.1
L058	44	Shrunk_Weight	26628.0	56628.0	SW: digit ‘2’ read as ‘5’
L068	40	Dressed_Weight	3979.4	33979.4	LW: decimal shift x0.1 + DW: digit ‘3’ inserted
L068	40	Total_Live_Weight	560080.0	56008.0	LW: decimal shift x0.1 + DW: digit ‘3’ inserted
L076	41	Shrunk_Weight	5286.2	52846.2	SW: digit ‘4’ inserted
L102	30	Shrunk_Weight	4827.8	48727.8	SW: digit ‘7’ inserted
L123	35	Shrunk_Weight	49954.4	59954.4	SW: digit ‘4’ read as ‘5’
L136	32	Shrunk_Weight	5185.3	52185.3	SW: digit ‘2’ inserted
L143	41	Shrunk_Weight	79760.0	59760.0	SW: digit ‘7’ read as ‘5’
L150	35	Dressed_Weight	3571.0	33571.0	DW: digit ‘3’ inserted
L151	43	Dressed_Weight	3478.8	31478.8	DW: digit ‘1’ inserted
L163	38	Shrunk_Weight	5759.6	57459.6	SW: digit ‘4’ inserted
L164	34	Dressed_Weight	356870.4	35687.0	DW: decimal shift x0.1
L170	36	Shrunk_Weight	54272.2	54722.2	SW: adjacent digits transposed
L172	34	Shrunk_Weight	46648.6	48648.6	SW: digit ‘6’ read as ‘8’
L182	40	Yield_pct	6563.0	65.6	scale Yield x0.01
L186	37	Yield_pct	6397.0	64.0	scale Yield x0.01
L195	34	Dressed_Weight	27197.4	28197.4	DW: digit ‘7’ read as ‘8’
L201	45	Shrunk_Weight	53547.8	56547.8	SW: digit ‘3’ read as ‘6’
L224	36	Dressed_Weight	4068.6	34068.6	DW: digit ‘3’ inserted
L235	35	Dressed_Weight	1751.2	31751.2	DW: digit ‘3’ inserted
L239	36	Shrunk_Weight	46361.4	49361.4	SW: digit ‘6’ read as ‘9’
L245	42	Shrunk_Weight	60132.6	61032.6	SW: adjacent digits transposed
L248	42	Shrunk_Weight	60132.6	61032.6	SW: adjacent digits transposed
L257	33	Shrunk_Weight	50826.6	51826.6	SW: digit ‘0’ read as ‘1’
L260	21	Yield_pct	6442.0	64.4	scale Yield x0.01
L262	43	Shrunk_Weight	49760.0	59760.0	SW: digit ‘4’ read as ‘5’
L271	19	Shrunk_Weight	280863.0	28086.0	SW: digit ‘3’ deleted

What happened to the flagged lots

outcome	n_values
corrected: DW	6
corrected: LW	3
corrected: SW	16
corrected: scale	3
deleted (no slip explains the value)	6
left for manual review	1
retained (sub-3-point yield discrepancy, weights plausible)	4

Lots the rules refused to touch - please check these against the source records

Received_Lot_nr	Head	Total_Live_Weight	Shrunk_Weight	Dressed_Weight	Yield_recorded	Yield_from_weights	rule
L244	38	48720	48132.8	37183.2	63.96	77.25	review: 13.3 pt yield discrepancy, no typing slip explains it

Lots usable in the weight / yield models, before and after correction

Outcome	Before	After	Gained
live weight	263	264	1
shrunk weight	260	264	4
dressed weight	258	262	4
yield	265	266	1

2.4 Derived variables and QC screen

QC screen after correction. `qc_weights` (all four weight checks at once) gates every weight / yield model.

Check	Observed	Failing
DD_count within [0, Head]	269	0
live wt/head 1000-2100 lb	265	0
shrunk wt/head 1000-2100 lb	266	0
dressed wt/head 600-1300 lb	264	0
yield 50-75%	268	0
recorded yield == Dressed/Shrunk (after repair)	262	2
YG panel sums > 0	229	0
Grade panel sums > 0	226	0

```
## 'data.frame':	455 obs. of  13 variables:
##  $ id             : Factor w/ 76 levels "P01","P02","P03",..: 28 28 28 28 37 38 51 51 51 51 ...
##  $ Producer_clean : chr  "P28" "P28" "P28" "P28" ...
##  $ Season         : Factor w/ 4 levels "2_Spring","1_Winter",..: 3 3 4 4 4 4 4 4 4 4 ...
##  $ Gender         : Factor w/ 2 levels "H","M": NA NA NA NA NA NA NA NA NA NA ...
##  $ Head           : num  38 38 38 38 23 45 32 33 33 32 ...
##  $ DD_count       : num  1 2 24 11 10 1 9 5 7 3 ...
##  $ prev           : num  0.0263 0.0526 0.6316 0.2895 0.4348 ...
##  $ dressed_wt_head: num  970 981 950 969 1024 ...
##  $ yield_pct_clean: num  63.6 63.1 64 63.5 61.7 ...
##  $ wt_corrected   : logi  FALSE FALSE FALSE FALSE FALSE FALSE ...
##  $ qc_weights     : logi  TRUE TRUE TRUE TRUE TRUE TRUE ...
##  $ yg_n           : num  38 38 38 38 23 45 NA NA NA NA ...
##  $ grade_n        : num  38 38 37 37 23 44 NA NA NA NA ...
```

3 Descriptive analysis

3.1 Structure of the data

Data structure

Item	Value
Lots	455
Producers (id)	76
Lots per producer (median [min-max])	3 [1-77]
Date received	2024-08-23 to 2025-09-11
Head per lot (median [IQR])	38 [35-41]
Lots with DD recorded	269
Lots with harvest weights	264
Lots with carcass grades	230

Missingness (n = all lots)

Variable	n_missing	pct_missing
Select	228	50.1
YG1	226	49.7
G137_Prime	226	49.7
YG2	225	49.5
YG3	225	49.5
YG4	225	49.5
YG5more	225	49.5
CAB_Prime	225	49.5
Cab_Upper_2_3rd	225	49.5
G137_Upper_2_third	225	49.5
Choice	225	49.5
Dressed_Weight	191	42.0

Variable	n_missing	pct_missing
Total_Live_Weight	190	41.8
Shrunk_Weight	189	41.5
Yield_pct	187	41.1
DD_count	186	40.9
prev	186	40.9
Gender	49	10.8
Season	2	0.4
Head	0	0.0

3.2 Numeric summaries

Lot-level descriptive statistics

Variable	n	missing	mean	sd	p25	median	p75	min	max
CAB_Prime	230	225	10.72	9.36	5.00	8.00	15.00	0.00	90.00
Cab_Upper_2_3rd	230	225	14.50	5.82	11.00	14.00	18.00	0.00	31.00
Choice	230	225	8.69	6.07	4.00	8.00	12.75	0.00	30.00
DD_count	269	186	3.65	4.43	1.00	2.00	4.00	0.00	27.00
G137_Prime	229	226	1.15	2.54	0.00	0.00	1.00	0.00	25.00
G137_Upper_2_third	230	225	2.46	2.98	0.00	1.00	4.00	0.00	15.00
Head	455	0	37.76	4.66	35.00	38.00	41.00	15.00	46.00
Select	227	228	0.28	1.01	0.00	0.00	0.00	0.00	10.00
YG1	229	226	1.01	1.43	0.00	0.00	2.00	0.00	8.00
YG2	230	225	6.49	5.32	2.00	5.00	9.00	0.00	25.00
YG3	230	225	18.95	4.93	15.00	19.00	22.00	7.00	31.00
YG4	230	225	9.94	5.70	6.00	9.00	14.00	0.00	25.00
YG5more	230	225	1.36	1.83	0.00	1.00	2.00	0.00	11.00
dressed_wt_head	264	191	928.95	91.69	866.71	924.38	983.28	688.81	1272.07
live_wt_head	265	190	1482.84	143.85	1385.00	1483.68	1570.00	1134.42	2041.82
prev_pct	269	186	9.63	11.30	2.86	5.88	11.11	0.00	65.85
shrunk_wt_head	266	189	1470.50	140.92	1375.82	1471.28	1562.38	1134.42	2021.40
yield_pct_clean	268	187	63.11	1.10	62.27	63.12	63.98	60.10	65.63

```
##
##
## Table: Season
##
## |Season   |   n|
## |:--------|---:|
## |2_Spring | 203|
## |1_Winter |   5|
## |3_Summer | 208|
## |4_Fall   |  37|
## |NA       |   2|
##
##
## Table: Gender
##
## |Gender |   n|
## |:------|---:|
## |H      |  84|
## |M      | 322|
## |NA     |  49|
##
##
## Table: Season x Gender
##
## |Season   |  H|   M| NA|
## |:--------|--:|---:|--:|
## |2_Spring | 32| 139| 32|
## |1_Winter |  0|   0|  5|
## |3_Summer | 35| 171|  2|
## |4_Fall   | 17|  10| 10|
## |NA       |  0|   2|  0|
```

3.3 Distribution of DD prevalence

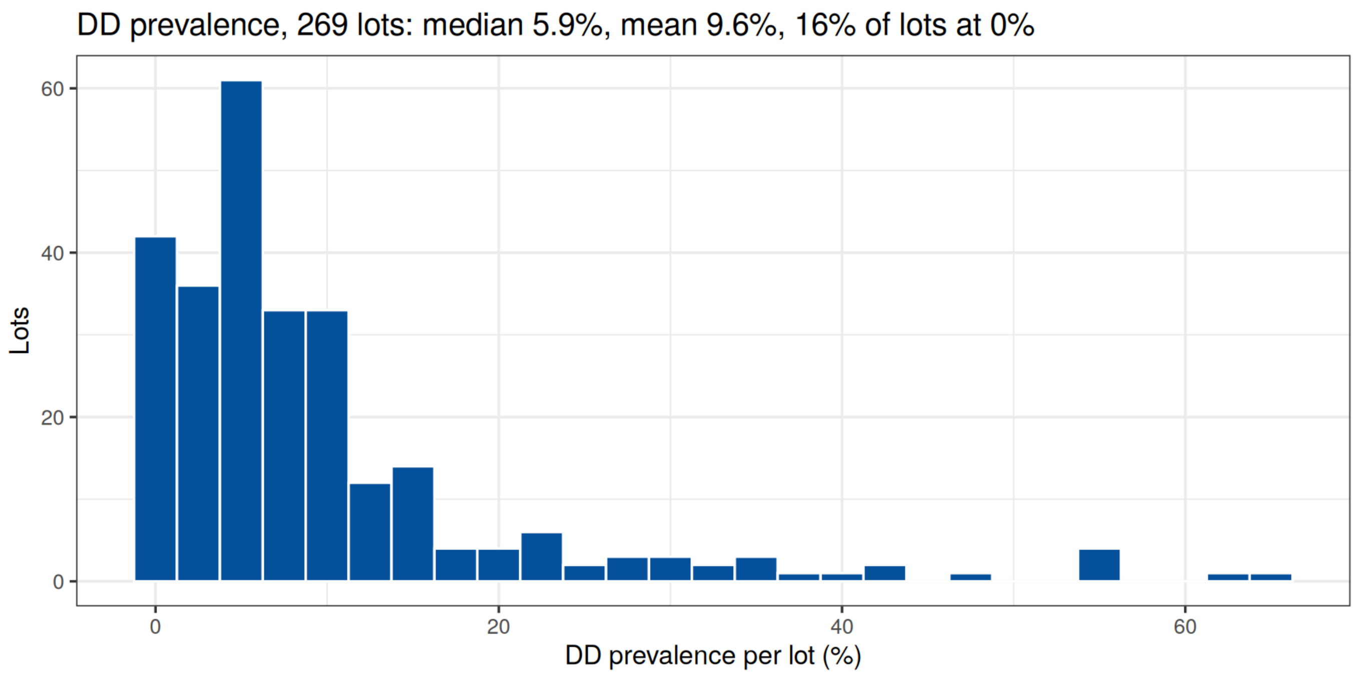

Pooled (animal-level) vs. average lot-level prevalence

lots	animals	dd_cases	pooled_prevalence_pct	mean_of_lot_prev_pct
269	10047	982	9.77	9.63

3.4 Producer-level view

Crude DD prevalence by producer (20 producers with most lots). sd_lot_pct = SD across the producer's lots; se_pooled_pct = binomial SE of the pooled rate; wilson_lo/hi = 95% Wilson interval.

id	Producer_clean	lots	head	dd	prev_pct	mean_lot_pct	sd_lot_pct	se_pooled_pct	sd_lo	sd_hi	wilson_lo	wilson_hi
P62	P62	45	1781	189	10.61	10.57	14.28	0.73	0.00	24.89	9.27	12.13
P68	P68	35	1494	129	8.63	8.59	5.37	0.73	3.26	14.01	7.31	10.17
P28	P28	22	781	93	11.91	11.39	13.70	1.16	0.00	25.61	9.82	14.37
P71	P71	20	806	214	26.55	26.31	17.77	1.56	8.78	44.32	23.62	29.71
P60	P60	13	485	9	1.86	1.87	3.56	0.61	0.00	5.41	0.98	3.49
P27	P27	9	303	35	11.55	11.99	12.51	1.84	0.00	24.06	8.42	15.64
P01	P01	8	289	18	6.23	5.76	3.05	1.42	3.17	9.28	3.98	9.63
P18	P18	8	256	18	7.03	7.03	7.23	1.60	0.00	14.26	4.49	10.84
P33	P33	7	266	12	4.51	4.51	3.94	1.27	0.57	8.45	2.60	7.72
P48	P48	7	220	18	8.18	8.21	3.69	1.85	4.49	11.88	5.24	12.56
P13	P13	6	209	8	3.83	3.84	2.36	1.33	1.47	6.19	1.95	7.37
P14	P14	6	210	22	10.48	10.48	7.81	2.11	2.67	18.28	7.02	15.35
P29	P29	5	177	10	5.65	5.75	3.51	1.74	2.14	9.16	3.10	10.09
P31	P31	5	185	16	8.65	8.65	5.86	2.07	2.79	14.51	5.39	13.59
P08	P08	4	148	3	2.03	2.08	2.66	1.16	0.00	4.69	0.69	5.79
P40	P40	4	143	13	9.09	9.11	2.75	2.40	6.35	11.84	5.39	14.93
P45	P45	4	144	4	2.78	2.60	3.00	1.37	0.00	5.78	1.09	6.92
P51	P51	4	130	24	18.46	18.47	8.05	3.40	10.41	26.51	12.73	26.00
P69	P69	4	136	6	4.41	4.41	2.94	1.76	1.47	7.35	2.04	9.29
P02	P02	3	108	5	4.63	4.63	5.78	2.02	0.00	10.41	1.99	10.38

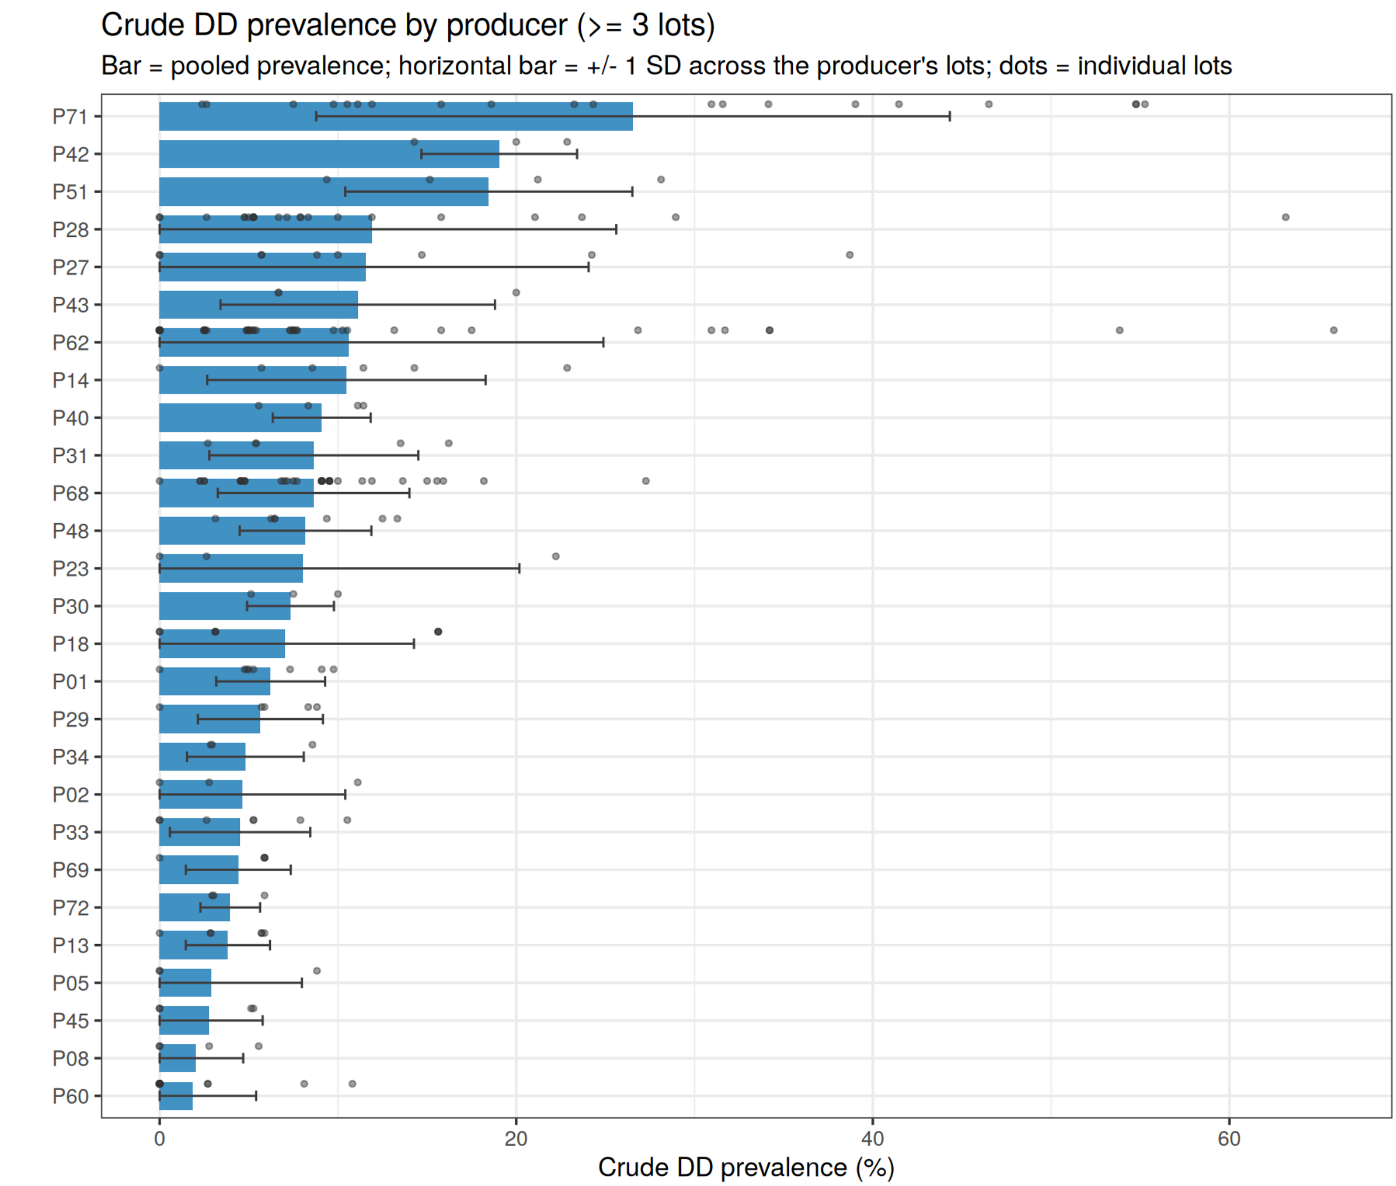
Crude DD prevalence by producer (>= 3 lots)
Bar = pooled prevalence; horizontal bar = +/- 1 SD across the producer's lots; dots = individual lots
P71
P42
P51
P28
P27
P43
P62
P14
P40
P31
P68
P48
P23
P30
P18
P01
P29
P34
P02
P33
P69
P72
P13
P05
P45
P08
P60
0
20
40
60
Crude DD prevalence (%)

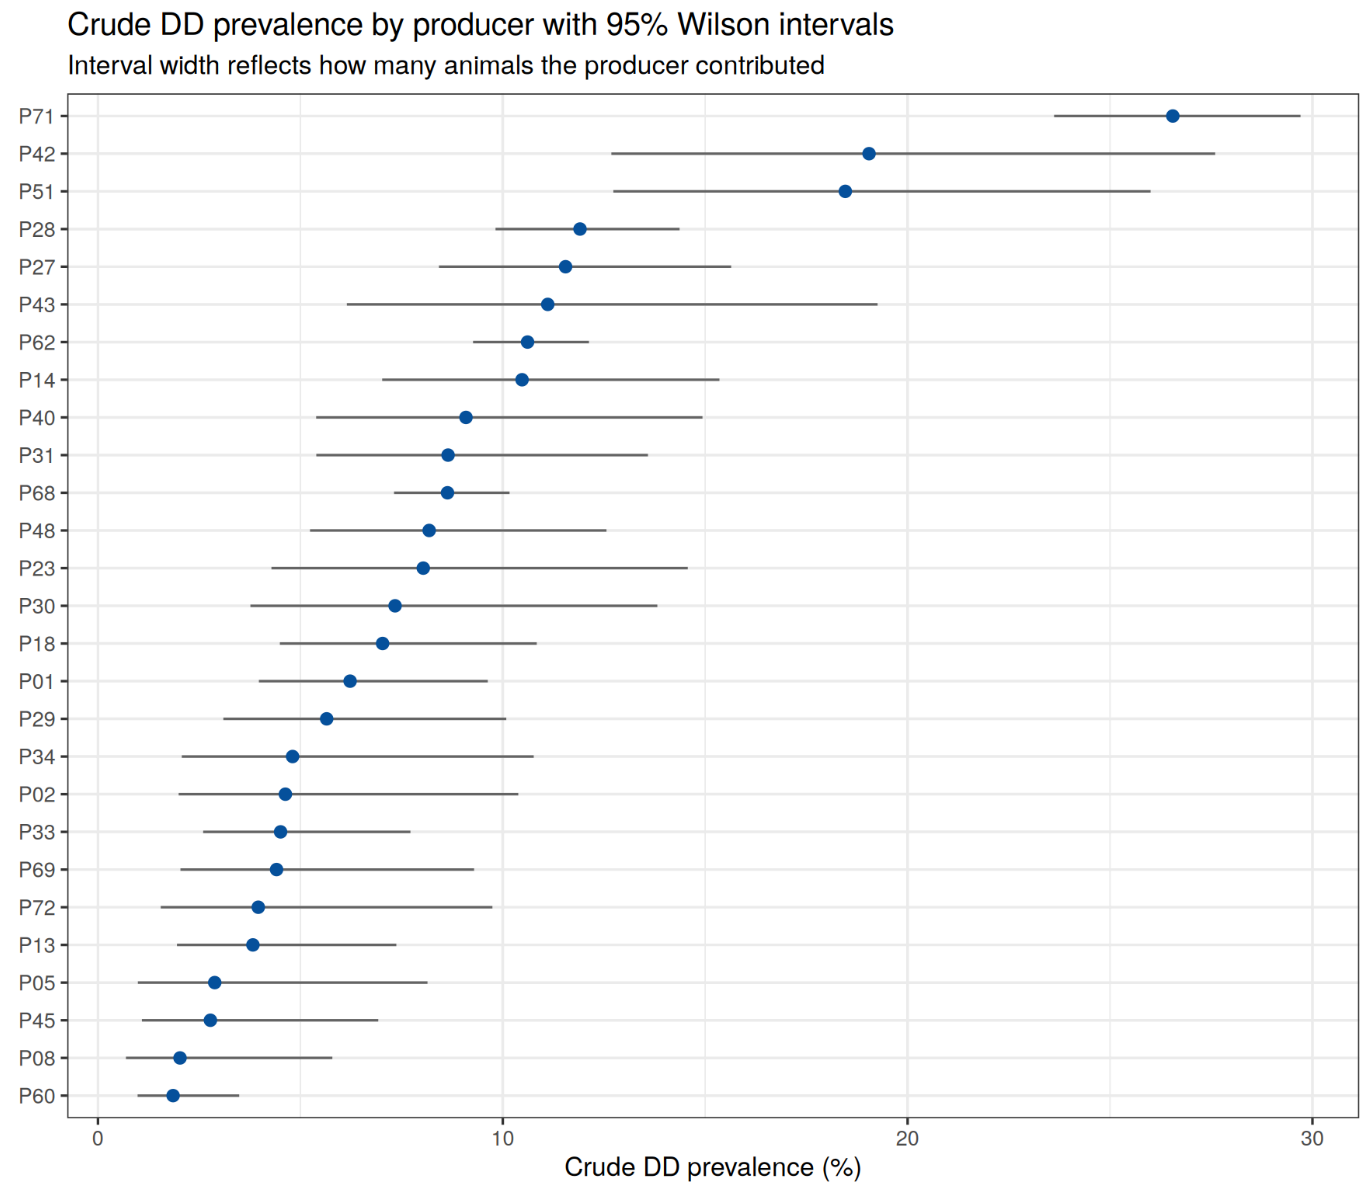

4 Bivariate analysis

4.1 Predictors of DD prevalence

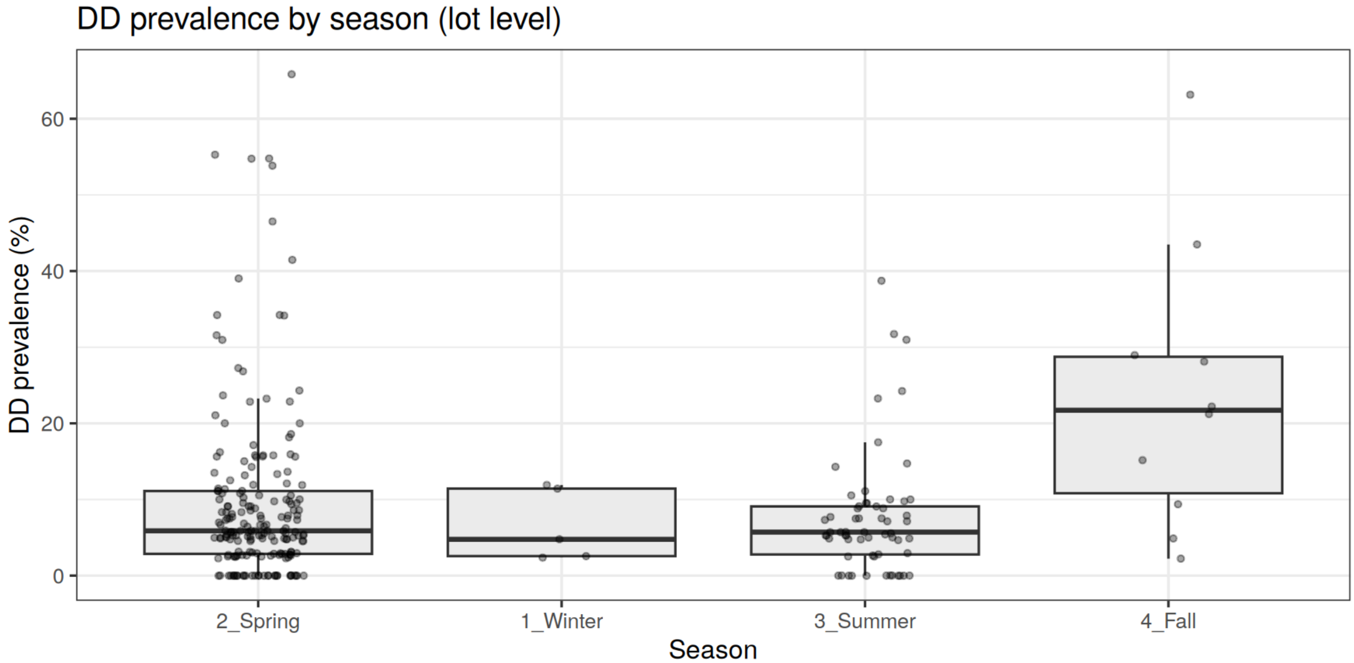

Unadjusted association with DD prevalence (lot level)

Predictor	Test	p
Season	Kruskal-Wallis	0.024900
Gender	Wilcoxon rank-sum	0.738000
Producer (id)	Kruskal-Wallis	0.000563
Head (lot size)	Spearman	0.227000

Crude binomial GLM (ignores producer clustering - for comparison with the GLMM only)

term	OR	CI	p

term	OR	CI	p
(Intercept)	0.134	0.11-0.16	0.00000
Season3_Summer	0.807	0.67-0.96	0.01830
GenderM	0.749	0.63-0.90	0.00153

4.2 DD prevalence vs. harvest outcomes (unadjusted)

Spearman correlation of lot DD prevalence with each harvest outcome (FDR across outcomes)

Outcome	n	rho	p	q_FDR
p_Select	216	0.17400	0.0106	0.153
p_G137_Prime	216	0.15800	0.0204	0.153
p_Cab_Upper_2_3rd	216	-0.14600	0.0321	0.161
p_G137_Upper_2_third	216	0.13300	0.0509	0.191
yield_pct_clean	256	0.08980	0.1520	0.456
p_Choice	216	0.06920	0.3110	0.717
p_YG3	219	0.06520	0.3370	0.717
p_YG1	219	0.05930	0.3820	0.717
p_CAB_Prime	216	0.01500	0.8270	0.936
dressed_wt_head	252	0.01230	0.8460	0.936
p_YG4	219	-0.01290	0.8500	0.936
p_YG2	219	-0.00863	0.8990	0.936
shrunk_wt_head	254	-0.00526	0.9340	0.936
live_wt_head	254	-0.00515	0.9350	0.936
p_YG5more	219	0.00541	0.9360	0.936

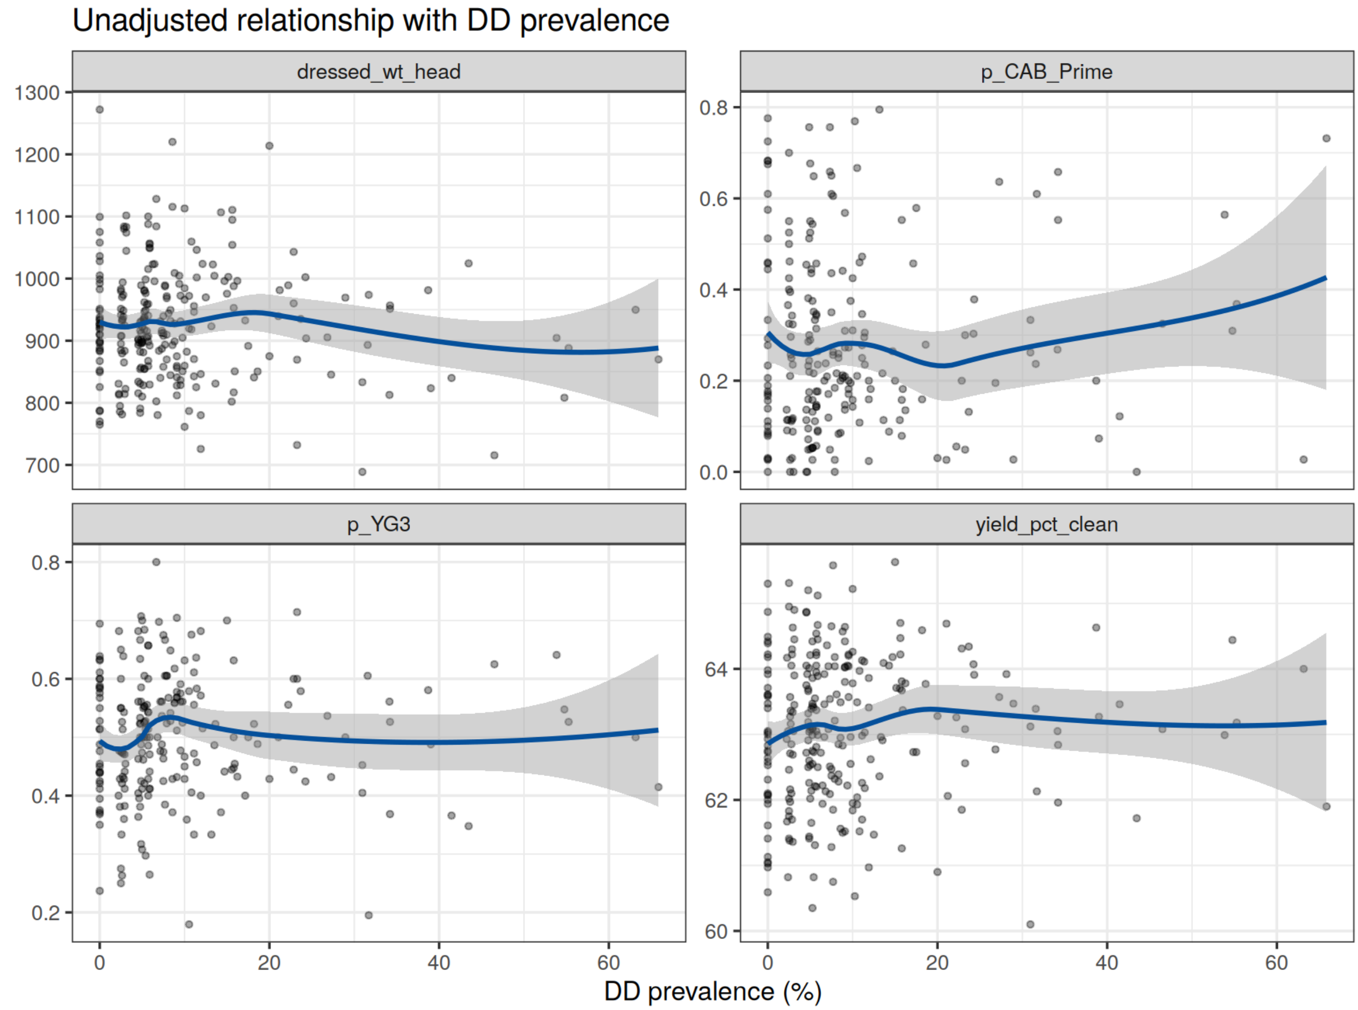

5 GLMM: DD prevalence with random producer

5.1 Primary model (Season, all lots)

```
## Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
##  Family: binomial  ( logit )
## Formula: cbind(DD_count, DD_neg) ~ Season + (1 | id)
##    Data: dat_prev
## Control: glmerControl(optimizer = "bobyqa", optCtrl = list(maxfun = 2e+05))
##
##      AIC      BIC   logLik deviance df.resid
##   1501.2   1519.2   -745.6   1491.2      264
##
## Scaled residuals:
##     Min      1Q  Median      3Q     Max
## -3.5571 -1.0450 -0.3457  0.5325 11.1418
##
## Random effects:
##  Groups Name        Variance Std.Dev.
##  id     (Intercept) 0.4671   0.6834
## Number of obs: 269, groups:  id, 48
##
## Fixed effects:
##                Estimate Std. Error z value Pr(>|z|)
## (Intercept)     -2.6921     0.1238 -21.748  < 2e-16 ***
## Season1_Winter  -0.3393     0.3042  -1.115  0.26466
## Season3_Summer  -0.2977     0.1097  -2.712  0.00668 **
## Season4_Fall     1.7881     0.2243   7.972 1.56e-15 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Correlation of Fixed Effects:
##             (Intr) Ssn1_W Ssn3_S
## Seasn1_Wntr -0.075
## Seasn3_Smmr -0.210  0.069
## Season4_Fll -0.257  0.065  0.117
```

```
##
## Lots in model: 269   Producers: 48   Seasons: 2_Spring, 1_Winter, 3_Summer, 4_Fall
```

```
##   chisq      df   ratio       p
## 768.342 264.000   2.910   0.000
```

```
## Producer variance = 0.467;  latent-scale ICC = 0.124
```

```
## AIC binomial GLMM = 1501.2 ; AIC beta-binomial GLMM = 1266.3 (lower is better)
##  Family: betabinomial  ( logit )
## Formula:          cbind(DD_count, DD_neg) ~ Season + (1 | id)
## Data: dat_prev
##
##      AIC      BIC   logLik deviance df.resid
##   1266.3   1287.9   -627.2   1254.3      263
##
## Random effects:
##
## Conditional model:
##  Groups Name        Variance Std.Dev.
##  id     (Intercept) 0.2195   0.4685
## Number of obs: 269, groups:  id, 48
##
## Dispersion parameter for betabinomial family (): 16.5
##
## Conditional model:
##                Estimate Std. Error z value Pr(>|z|)
## (Intercept)     -2.4718     0.1229 -20.115  < 2e-16 ***
## Season1_Winter  -0.1659     0.4320  -0.384    0.701
## Season3_Summer  -0.1985     0.1620  -1.226    0.220
## Season4_Fall     1.3159     0.3073   4.283 1.85e-05 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
##
## Table: Beta-binomial GLMM: DD prevalence by season (preferred inference)
##
## |Season   |   prob|     SE|  df| asymp.LCL| asymp.UCL|
## |:--------|------:|------:|---:|---------:|---------:|
## |2_Spring | 0.0779| 0.0088| Inf|    0.0622|    0.0970|
## |1_Winter | 0.0668| 0.0273| Inf|    0.0294|    0.1444|
## |3_Summer | 0.0647| 0.0106| Inf|    0.0469|    0.0888|
## |4_Fall   | 0.2394| 0.0522| Inf|    0.1521|    0.3558|
##
##
## Table: Beta-binomial GLMM: pairwise season contrasts
##
## |contrast            | odds.ratio|     SE|  df| null| z.ratio| p.value|
## |:-------------------|----------:|------:|---:|----:|-------:|-------:|
## |2_Spring / 1_Winter |     1.1804| 0.5099| Inf|    1|  0.3840|  0.9807|
## |2_Spring / 3_Summer |     1.2196| 0.1976| Inf|    1|  1.2257|  0.6104|
## |2_Spring / 4_Fall   |     0.2682| 0.0824| Inf|    1| -4.2826|  0.0001|
## |1_Winter / 3_Summer |     1.0332| 0.4624| Inf|    1|  0.0730|  0.9999|
## |1_Winter / 4_Fall   |     0.2272| 0.1177| Inf|    1| -2.8615|  0.0219|
## |3_Summer / 4_Fall   |     0.2199| 0.0726| Inf|    1| -4.5852|  0.0000|
```

Odds ratios for DD, reference season = Spring

term	OR	CI95	p
(Intercept)	0.068	0.05-0.09	0.00000
Season1_Winter	0.712	0.39-1.29	0.26500
Season3_Summer	0.743	0.60-0.92	0.00668
Season4_Fall	5.978	3.85-9.28	0.00000

5.2 Estimated marginal means and contrasts

Model-based DD prevalence by season (probability scale, producer-averaged)

Season	prob	SE	df	asymp.LCL	asymp.UCL
2_Spring	0.0634	0.0074	Inf	0.0505	0.0795
1_Winter	0.0460	0.0140	Inf	0.0251	0.0828
3_Summer	0.0479	0.0067	Inf	0.0363	0.0629
4_Fall	0.2882	0.0465	Inf	0.2062	0.3871

Observed DD prevalence by season: pooled rate, mean of lot prevalences, SD and SE across lots

Season	lots	pooled_pct	mean_lot_pct	sd_lot_pct	se_lot_pct
2_Spring	193	9.87	9.61	11.39	0.82
1_Winter	5	6.50	6.61	4.71	2.11
3_Summer	61	7.75	7.60	7.93	1.02
4_Fall	10	22.79	23.88	18.55	5.87

Pairwise season contrasts (Tukey-adjusted, odds-ratio scale)

contrast	odds.ratio	SE	df	null	z.ratio	p.value
2_Spring / 1_Winter	1.4040	0.4271	Inf	1	1.1154	0.6800
2_Spring / 3_Summer	1.3467	0.1478	Inf	1	2.7123	0.0338
2_Spring / 4_Fall	0.1673	0.0375	Inf	1	-7.9723	0.0000
1_Winter / 3_Summer	0.9592	0.3033	Inf	1	-0.1318	0.9992
1_Winter / 4_Fall	0.1191	0.0436	Inf	1	-5.8107	0.0000
3_Summer / 4_Fall	0.1242	0.0295	Inf	1	-8.7694	0.0000

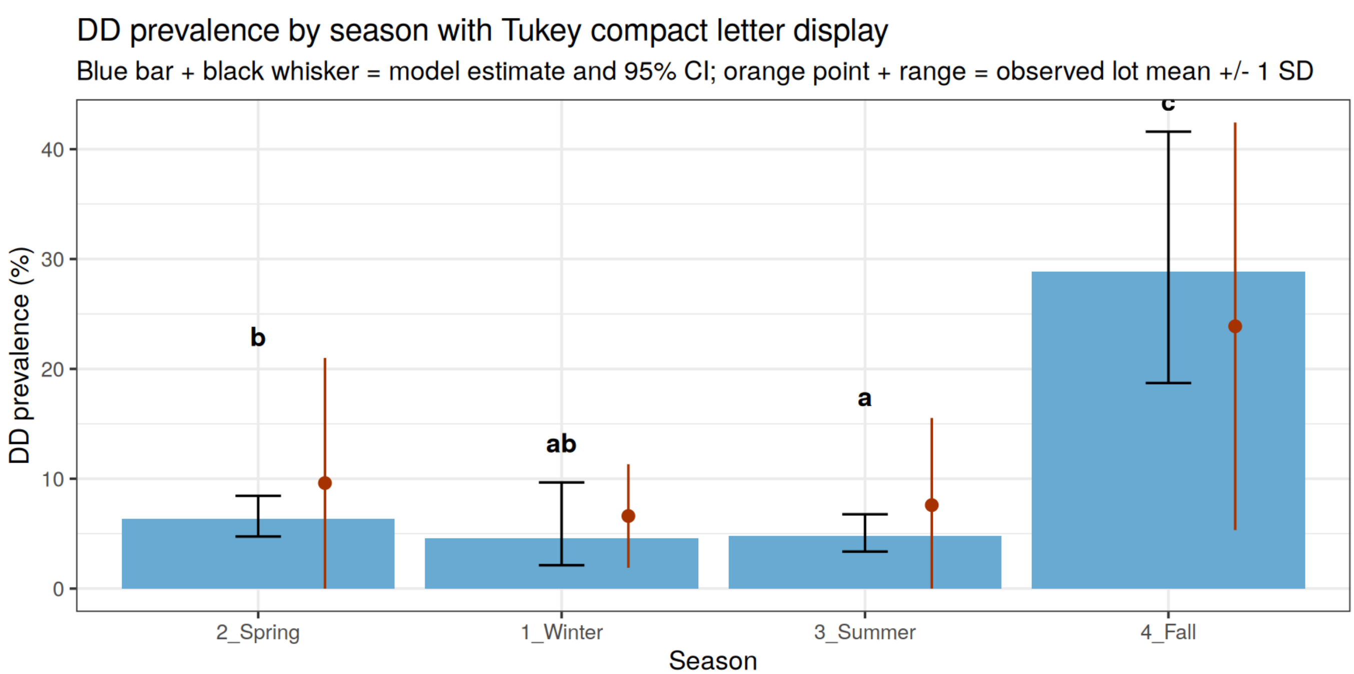

5.3 Producer random effects

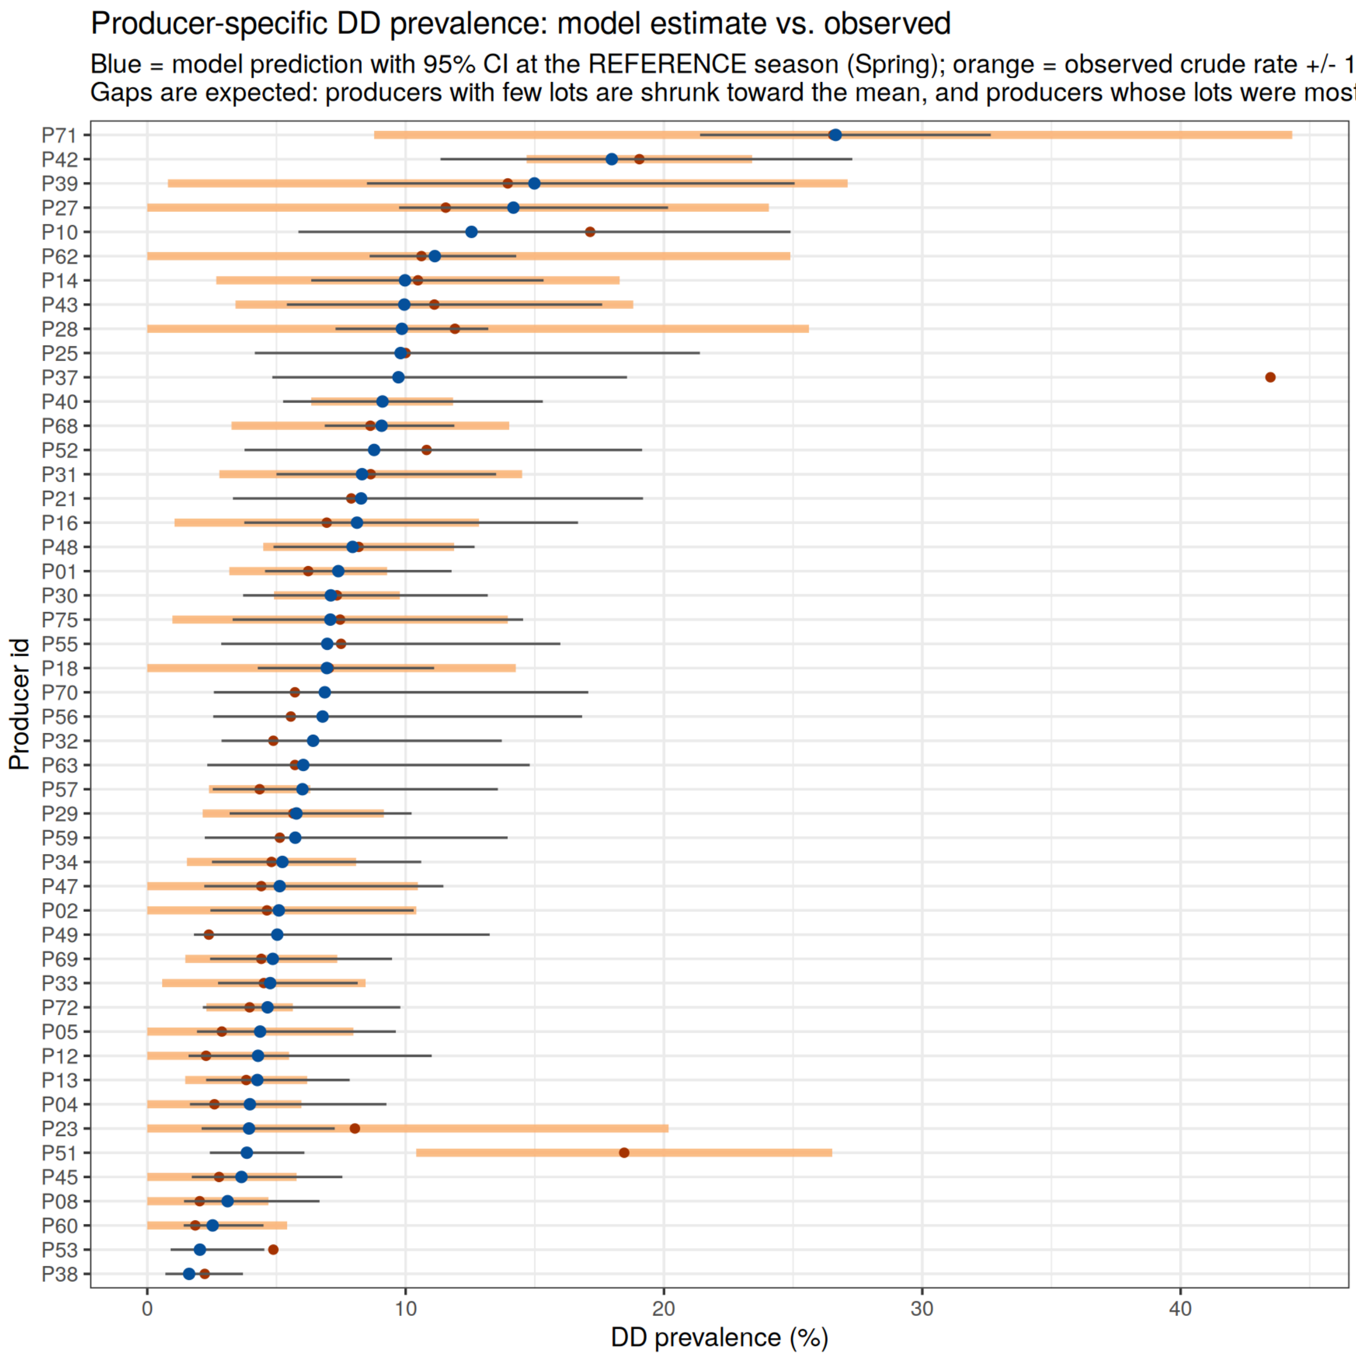

Ten producers with the highest model-predicted DD prevalence, with observed rate and lot-to-lot SD

id	Producer_clean	lots	predicted_pct	CI	observed_pct	sd_across_lots
P71	P71	20	26.6	21.4-32.7	26.6	17.8
P42	P42	3	18.0	11.3-27.3	19.0	4.4
P39	P39	2	15.0	8.5-25.1	14.0	13.2
P27	P27	9	14.2	9.7-20.2	11.6	12.5
P10	P10	1	12.6	5.8-24.9	17.1	NA
P62	P62	45	11.1	8.6-14.3	10.6	14.3
P14	P14	6	10.0	6.3-15.3	10.5	7.8

id	Producer_clean	lots	predicted_pct	CI	observed_pct	sd_across_lots
P43	P43	3	9.9	5.4-17.6	11.1	7.7
P28	P28	22	9.9	7.3-13.2	11.9	13.7
P25	P25	1	9.8	4.2-21.4	10.0	NA

5.4 Gender-adjusted model (Winter not estimable)

```
## Gender-adjusted frame: n = 220 lots, seasons present = 2_Spring, 3_Summer
```

```
## Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
##  Family: binomial  ( logit )
## Formula: cbind(DD_count, DD_neg) ~ Season + Gender + (1 | id)
##    Data: dat_prev_g
## Control: glmerControl(optimizer = "bobyqa", optCtrl = list(maxfun = 2e+05))
##
##      AIC      BIC   logLik deviance df.resid
##   1166.7   1180.3   -579.4   1158.7      216
##
## Scaled residuals:
##     Min      1Q  Median      3Q     Max
## -3.2753 -1.0398 -0.2825  0.4651  8.7132
##
## Random effects:
##  Groups Name        Variance Std.Dev.
##  id     (Intercept) 0.392    0.6261
## Number of obs: 220, groups:  id, 39
##
## Fixed effects:
##                Estimate Std. Error z value Pr(>|z|)
## (Intercept)     -2.8487     0.1555 -18.315   <2e-16 ***
## Season3_Summer  -0.1569     0.1148  -1.367   0.1717
## GenderM          0.2466     0.1130   2.182   0.0291 *
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Correlation of Fixed Effects:
##             (Intr) Ssn3_S
## Seasn3_Smmr -0.230
## GenderM     -0.589  0.042
```

Season + Gender adjusted odds of DD (Winter absent: all 5 Winter lots lack Gender)

term	OR	CI95	p
(Intercept)	0.058	0.04-0.08	0.0000
Season3_Summer	0.855	0.68-1.07	0.1720
GenderM	1.280	1.03-1.60	0.0291

Predicted DD prevalence by gender

Gender	prob	SE	df	asymp.LCL	asymp.UCL
H	0.0508	0.0074	Inf	0.0382	0.0674
M	0.0641	0.0076	Inf	0.0508	0.0807

6 Predicting harvest outcomes from DD prevalence

Model for every outcome:

```
outcome ~ prev + Season + Gender + (1 | id)
```

- continuous outcomes (per head): `lmer` , effect = change in the outcome per **+10 pp** DD prevalence;
- count outcomes: beta-binomial GLMM on `cbind(k, n - k)` (set `params$betabinomial = FALSE` for a plain binomial `glmer`), effect = **odds ratio** per +10 pp DD prevalence.

Prevalence is mean-centred inside each model, so the intercept and slope are on comparable scales; the slope, its CI and its p-value are unaffected.

Because Winter carries no Gender information, these models run on the Gender-complete frame (Spring / Summer / Fall). A Season-only sensitivity run that keeps Winter is included.

Effect of +10 percentage points DD prevalence on each harvest outcome, adjusted for Season and Gender, random producer. Gaussian outcomes: absolute change; count outcomes: odds ratio.

Outcome	Scale	Unit	n	producers	SD(producer)	Effect	p	q_FDR	note
yield_pct_clean	difference	yield % points	210	39	0.869	0.174 (0.061, 0.287)	0.00247	0.037	
Select	odds ratio	OR (share of graded carcasses)	181	35	0.527	1.406 (1.019, 1.940)	0.03800	0.267	
YG2	odds ratio	OR (share of graded carcasses)	184	35	0.734	0.878 (0.769, 1.002)	0.05350	0.267	
YG1	odds ratio	OR (share of graded carcasses)	184	35	0.632	0.829 (0.668, 1.029)	0.08860	0.332	
G137_Upper_2_third	odds ratio	OR (share of graded carcasses)	181	35	0.857	1.131 (0.960, 1.331)	0.14000	0.420	
YG5more	odds ratio	OR (share of graded carcasses)	184	35	0.610	1.115 (0.940, 1.323)	0.21000	0.524	
YG3	odds ratio	OR (share of graded carcasses)	184	35	0.219	1.028 (0.968, 1.092)	0.36600	0.755	
CAB_Prime	odds ratio	OR (share of graded carcasses)	181	35	0.681	0.958 (0.867, 1.059)	0.40300	0.755	
Cab_Upper_2_3rd	odds ratio	OR (share of graded carcasses)	181	35	0.130	0.971 (0.890, 1.058)	0.49900	0.755	
live_wt_head	difference	lb/head	208	39	94.910	-3.949 (-16.434, 8.536)	0.53500	0.755	
Choice	odds ratio	OR (share of graded carcasses)	181	35	0.523	1.035 (0.923, 1.161)	0.55400	0.755	
YG4	odds ratio	OR (share of graded carcasses)	184	35	0.543	1.019 (0.927, 1.121)	0.69300	0.791	
dressed_wt_head	difference	lb/head	207	39	60.766	1.303 (-6.914, 9.519)	0.75600	0.791	
shrunk_wt_head	difference	lb/head	209	39	92.214	-1.685 (-13.780, 10.411)	0.78500	0.791	
G137_Prime	odds ratio	OR (share of graded carcasses)	181	35	0.520	1.029 (0.832, 1.274)	0.79100	0.791	

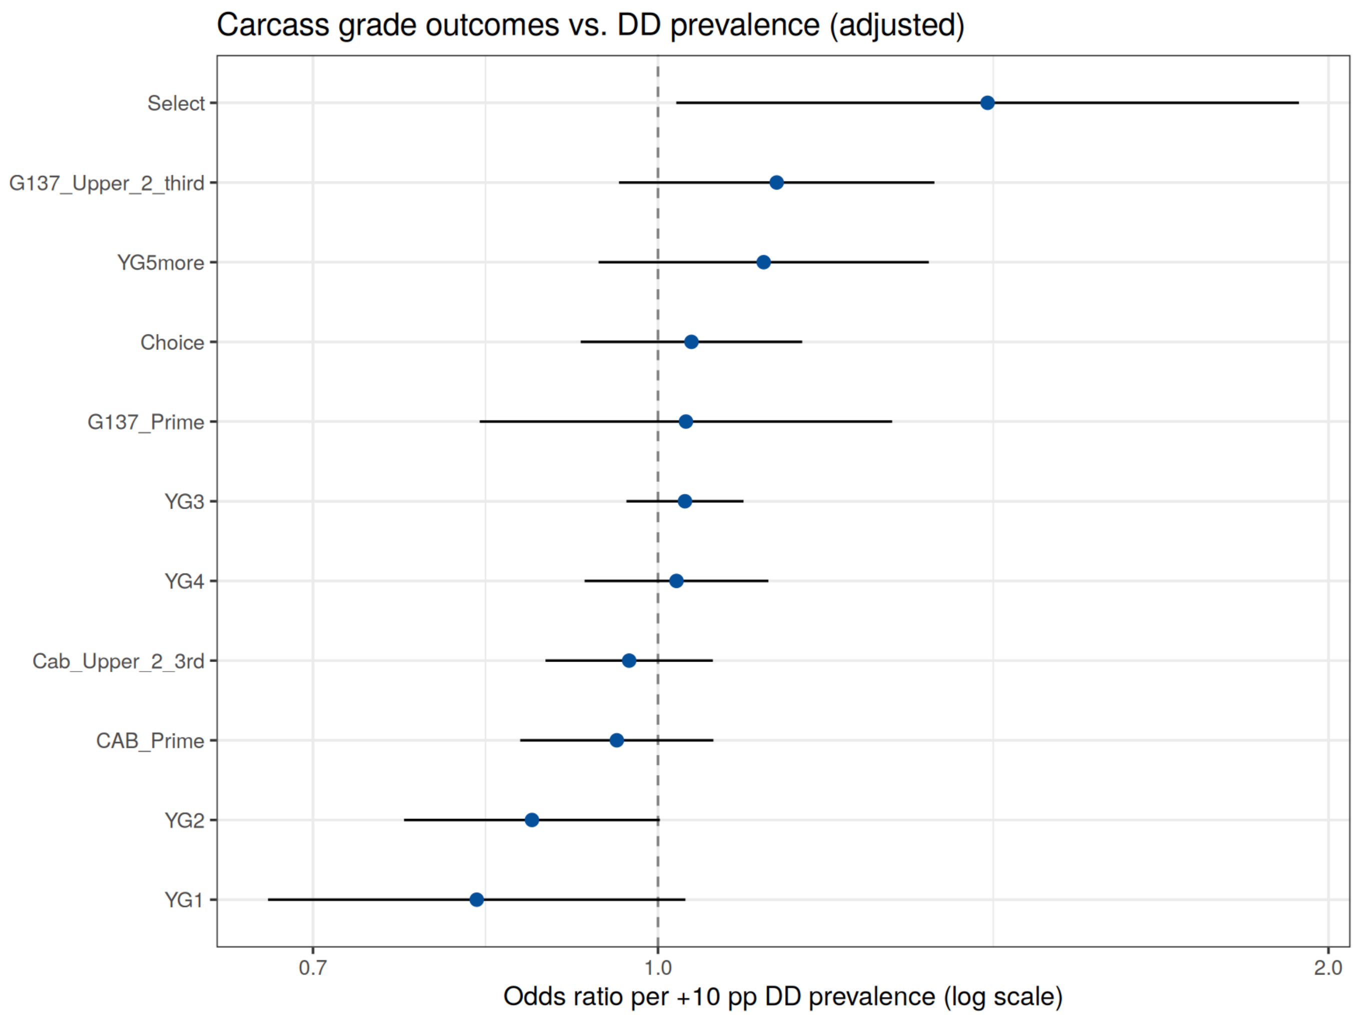
Carcass grade outcomes vs. DD prevalence (adjusted)
Select
G137_Upper_2_third
YG5more
Choice
G137_Prime
YG3
YG4
Cab_Upper_2_3rd
CAB_Prime
YG2
YG1
0.7
1.0
2.0
Odds ratio per +10 pp DD prevalence (log scale)

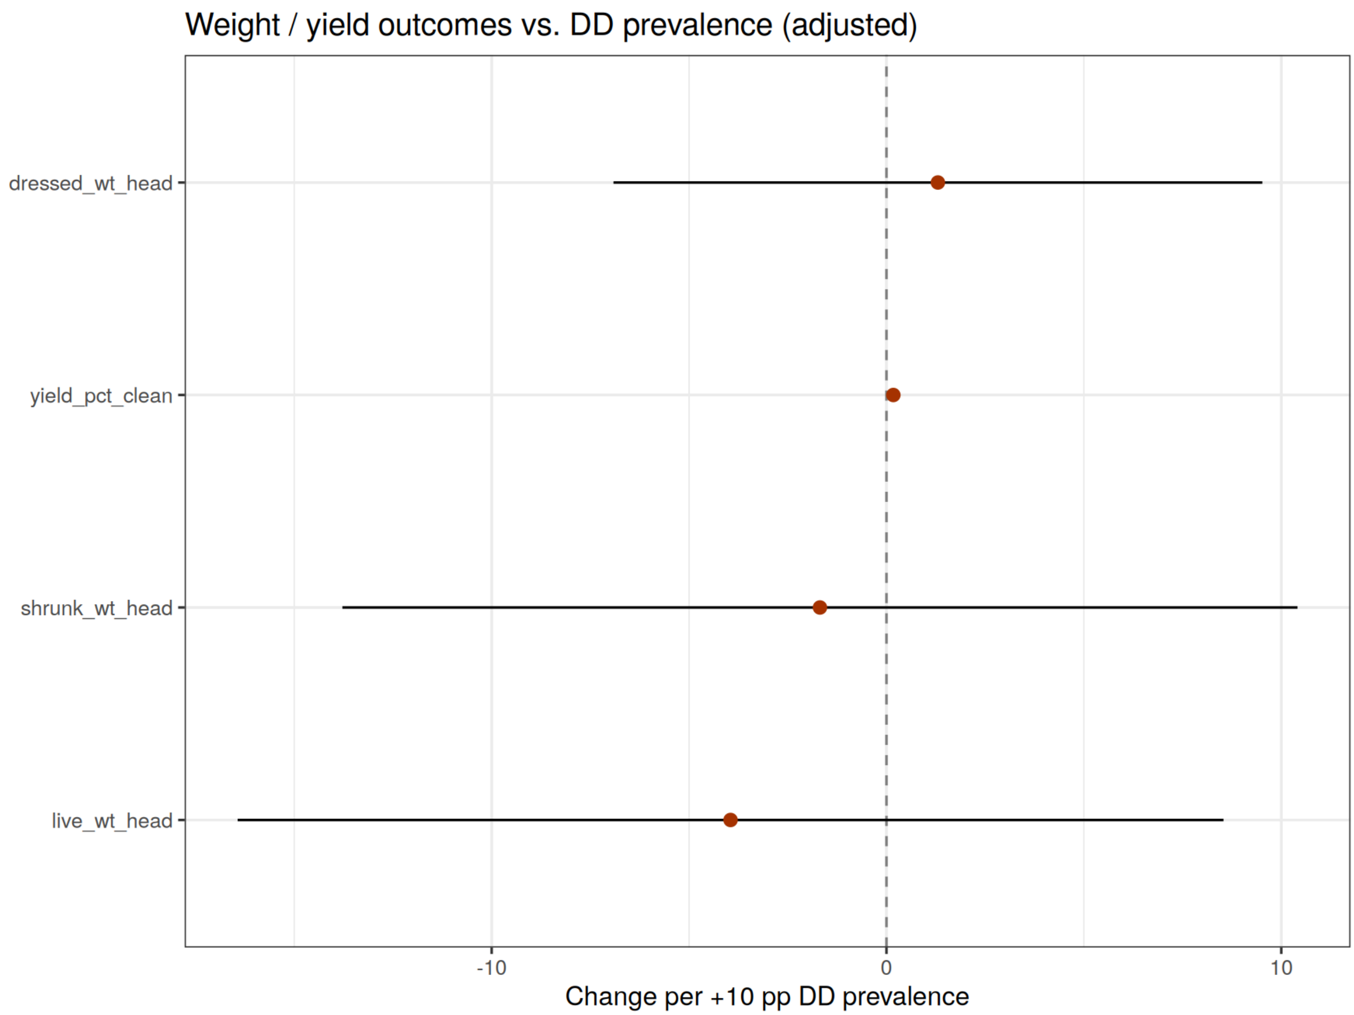

6.1 Why does yield move when neither weight does?

`Yield_pct` is a ratio, `100 x Dressed / Shrunk`. A ratio can move because its numerator rises, because its denominator falls, or because of how the denominator was produced. All three are checked below before the association is taken at face value.

Shrunk_Weight / Total_Live_Weight takes only 14 distinct values in 205 lots: it is a pencil shrink applied to the live weight, not a second weighing

shrink_grp	lots	mean DD prevalence (%)
0% pencil shrink	78	10.66
1% pencil shrink	118	7.82
3%+ pencil shrink	9	4.09

Decomposing the yield association

Model	Unit	Effect per +10 pp DD	t
Yield (Dressed/Shrunk), as reported	yield % points	0.187	3.115
Yield, additionally adjusted for the pencil-shrink band	yield % points	0.191	3.135
Dressing % on LIVE weight (no pencil shrink involved)	% points	0.229	3.685
Live weight per head	lb/head	-4.238	-0.657
Dressed weight per head	lb/head	0.558	0.129

At the mean lot (1472 lb live, 922 lb carcass per head), a +10 pp rise in DD prevalence is accompanied by **-4.24 lb of live weight** and **+0.56 lb of carcass**. Feeding those through the ratio gives +0.038 points from the carcass and +0.180 points from the live weight - i.e. **83% of the yield effect is the denominator falling, not the numerator rising**.

The ratio is also far better measured than its parts: yield has a coefficient of variation of 1.82% across lots, against 8.8% for live weight and 9.2% for carcass weight. Lot-to-lot differences in frame and finish move both weights together and cancel in the ratio, which is why a 0.19-point signal is detectable in yield while 4 lb and 1 lb are lost in the noise of the components.

What this rules in and out. The pencil-shrink band explains none of it: adjusting for it leaves the estimate unchanged, and computing dressing percentage straight off the live weight — where no contract shrink appears at all — gives a *larger* effect, not a smaller one. So the association is in the true carcass-to-live ratio. What it is *not* is extra saleable beef: dressed weight per head is flat (+0.6 lb per +10 pp, t = 0.13). A higher dressing percentage obtained by losing live weight is shrink, not gain — it is favourable only to whoever buys on the rail.

6.2 Could the yield result be confounded by days on feed?

The obvious rival explanation: cattle that stay on feed longer accumulate more exposure to DD infection pressure **and** arrive heavier and more finished, and finish raises dressing percentage. That would produce a positive DD-yield association with no causal link at all.

The honest starting point is that this data set contains no days-on-feed, age, arrival-weight or placement-date field. The hypothesis therefore cannot be tested directly, only probed with the finish measures that *are* recorded — mean yield grade, the YG4+ share, the marbling share, and live and carcass weight per head. Every model below already holds season, sex and producer constant, so the confounder would have to operate *within* a producer, across their own consecutive lots.

Step 1 - do the finish proxies actually rise with DD prevalence? (n = 175 lots, 34 producers)

Finish / size proxy	Change per +10 pp DD	SE	t	Direction the confounding story needs
Live weight per head (lb)	-4.13000	7.0500	-0.586	higher (longer on feed)
Carcass weight per head (lb)	0.76200	4.7200	0.162	higher (longer on feed)
Mean yield grade of the lot	0.03310	0.0234	1.410	higher (longer on feed)
Share grading YG4 or higher	0.00793	0.0114	0.693	higher (longer on feed)
Share in the upper two-thirds or better for marbling	-0.00950	0.0102	-0.927	higher (longer on feed)

Only one proxy leans the way the hypothesis requires, and weakly: **mean yield grade rises by 0.03 per +10 pp DD (t = 1.4)**, with fewer YG2 carcasses and marginally more YG5+. Live weight does **not** rise — the point estimate is negative. That is not decisive evidence against the hypothesis, because cattle are marketed to a target endpoint: out-weight is held roughly constant by design, which destroys slaughter weight as a proxy for days on feed. Carcass finish is the better proxy, and it moves only slightly.

Step 2 - does adjusting for finish and size take the DD effect away?

Additionally adjusted for	DD effect on yield per +10 pp	t	% of the unadjusted effect remaining
none	0.200	3.125	100.000
meanYG	0.178	2.872	88.914
fat_hi	0.192	3.066	95.842
marb_hi	0.197	3.064	98.499
live weight	0.207	3.281	103.468
carcass weight	0.196	3.279	98.001
meanYG + marbling + live weight	0.170	2.739	85.135

How much confounding would it take? One yield-grade unit is worth 0.71 yield points (t = 3.5), and +10 pp of DD buys 0.033 of a yield grade. Multiply the two and the whole yield-grade pathway carries **0.023 of the 0.200-point effect - about 12%**. For days on feed to explain the rest, the unmeasured part of it would have to be roughly **9 times more strongly linked to both DD and dressing percentage than mean yield grade is** - a large ask for a variable whose visible footprint on this data set is so faint.

Two further checks point the same way. Holding live weight fixed explicitly, a log-log model gives +0.40% of carcass weight per +10 pp DD (t = 3.91) with a live-weight exponent of 1.05, so the effect is not an artefact of lots differing in size. And restricting to the 120 lots in a narrow 1,350-1,600 lb live-weight band - crude matching on finished weight - leaves +0.186 points per +10 pp (t = 2.30), close to the full-sample estimate.

Verdict. Days on feed is a live hypothesis and this data set cannot rule it out — there is no variable in the file that measures it. What the data can say is narrower and worth stating exactly: the *measurable* finish signal accounts for roughly an eighth of the yield association, the effect survives explicit adjustment for size and crude matching on slaughter weight, and the one proxy that moves in the predicted direction does so weakly. Anyone reporting the yield result should carry this limitation with it.

What would settle it. Three fields, none of them hard to capture at receipt: **placement date or days on feed**, **arrival (in-)weight**, and **source or background group**. With days on feed recorded, the question becomes a single covariate rather than an argument.

Sensitivity: all lots (corrections applied) vs. only lots that never needed a correction. Similar effects mean the repairs are not driving the result.

Outcome	all_effect	all_p	all_n	unrep_effect	unrep_p	unrep_n
live_wt_head	-3.950	0.53500	208	-7.880	0.23800	179
shrunk_wt_head	-1.680	0.78500	209	-7.240	0.27800	179
dressed_wt_head	1.300	0.75600	207	-1.630	0.71600	178
yield_pct_clean	0.174	0.00247	210	0.194	0.00246	179
YG1	0.829	0.08860	184	0.846	0.19400	155
YG2	0.878	0.05350	184	0.868	0.05590	155
YG3	1.030	0.36600	184	1.030	0.43700	155
YG4	1.020	0.69300	184	1.030	0.61800	155
YG5more	1.120	0.21000	184	1.090	0.39000	155
CAB_Prime	0.958	0.40300	181	0.982	0.73600	153
G137_Prime	1.030	0.79100	181	1.090	0.44600	153
Cab_Upper_2_3rd	0.971	0.49900	181	0.963	0.39600	153
G137_Upper_2_third	1.130	0.14000	181	1.110	0.28200	153
Choice	1.040	0.55400	181	0.989	0.86400	153
Select	1.410	0.03800	181	1.400	0.03450	153

Sensitivity: Season+Gender adjusted vs. Season-only (Winter retained)

Outcome	adj_effect	adj_p	sens_effect	sens_p
live_wt_head	-3.950	0.53500	-0.161	0.97800
shrunk_wt_head	-1.680	0.78500	2.280	0.69400
dressed_wt_head	1.300	0.75600	2.860	0.46300
yield_pct_clean	0.174	0.00247	0.135	0.00466
YG1	0.829	0.08860	0.957	0.61500
YG2	0.878	0.05350	0.894	0.05800
YG3	1.030	0.36600	1.020	0.46700
YG4	1.020	0.69300	1.010	0.77900
YG5more	1.120	0.21000	1.090	0.22500
CAB_Prime	0.958	0.40300	1.010	0.88200

Outcome	adj_effect	adj_p	sens_effect	sens_p
G137_Prime	1.030	0.79100	1.010	0.94400
Cab_Upper_2_3rd	0.971	0.49900	0.940	0.11000
G137_Upper_2_third	1.130	0.14000	1.110	0.17100
Choice	1.040	0.55400	1.040	0.42500
Select	1.410	0.03800	1.320	0.01780

Predicted dressed weight per head (lb) at fixed DD prevalence, averaged over Season and Gender

DD_prevalence_pct	emmean	SE	lower.CL	upper.CL
0	903.3	12.0	879.2	927.4
5	903.9	11.7	880.4	927.5
10	904.6	11.7	881.0	928.2
20	905.9	12.9	880.1	931.7
30	907.2	15.2	877.1	937.3

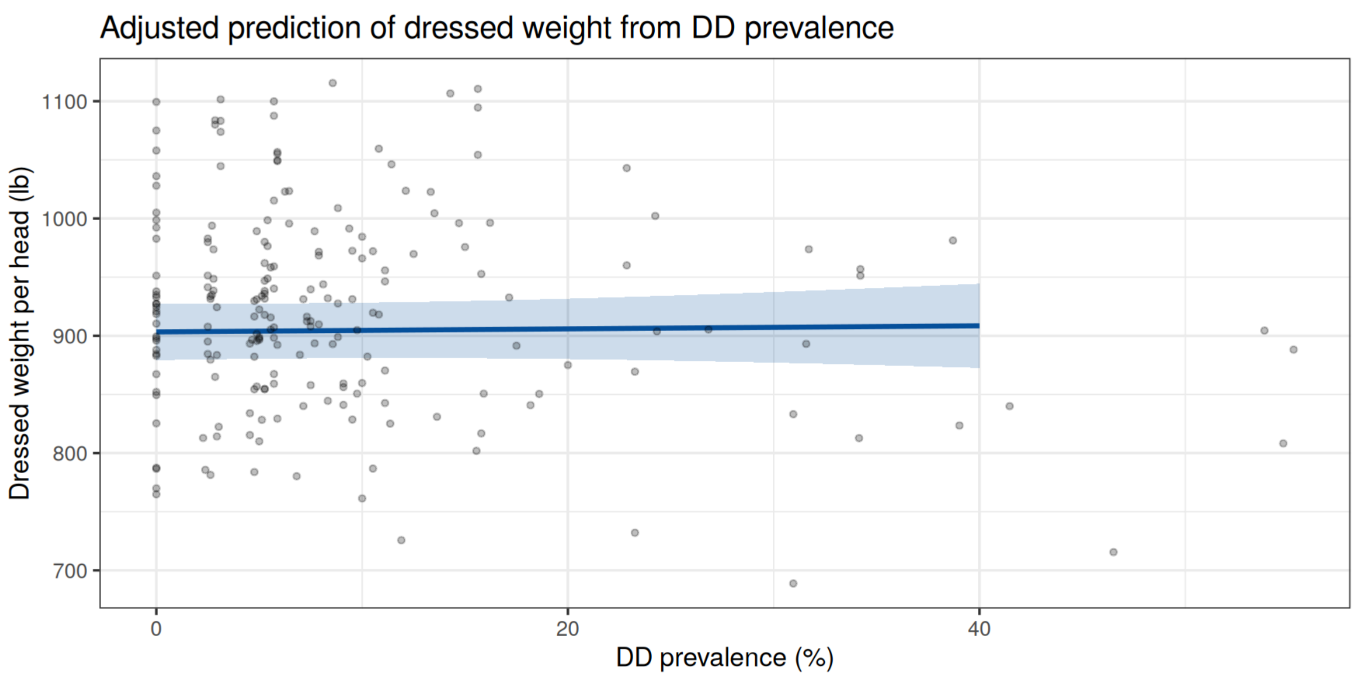

7 Summary of findings

- **Data recovered.** 76 producers after name cleaning (from 93 raw factor levels); the prevalence model now uses **269 lots / 48 producers** instead of the 49 lots the old `Origin` variable allowed.
- **DD burden.** 982 DD cases in 10,047 animals, pooled prevalence **9.8%**; 16% of lots have no DD at all.
- **Producer matters.** Producer variance 0.47 on the logit scale (ICC 0.12), i.e. about 12% of the variation in DD risk sits between producers - the random `id` term is doing real work.
- **Season.** Model-based prevalence: Spring 6.3%, Summer 4.8%, Fall 28.8%, Winter 4.6% (5 Feb lots only). Fall is far above every other season (all Tukey $p < 0.001$).
- **Data corrected.** 28 mistyped values in 27 lots repaired (each explained by one keystroke slip), 6 values deleted as unreconstructable, 1 lot left for manual review; that returns 262 lots to the weight/yield models.
- **Overdispersion.** Pearson chi2/df = 2.91; the beta-binomial GLMM fits far better (AIC 1266 vs 1501), so its wider intervals are the honest ones.
- **Harvest outcomes.** After adjustment for Season, Gender and producer, and FDR correction across the 15 outcomes, the surviving association is yield_pct_clean (0.174, q = 0.037).

Interpretation caveats worth stating in any write-up: DD status and harvest data are recorded on the same lot at the same time, so these are **associations, not causal effects**; DD was recorded for 269 of 455 lots and carcass grades for 230, and that missingness is not random across producers; and Fall/Winter rest on 37 and 5 lots respectively, so the seasonal contrast needs a second year of data before it is taken as a stable seasonal pattern rather than a property of these particular autumn consignments.

8 Verification

Automated checks

Check	Pass
id has no missing values	TRUE
every lot mapped to a producer	TRUE
id count < raw Producer levels	TRUE
no lot has DD_count > Head	TRUE
Season has all four levels defined	TRUE
Winter lots retained in the data	TRUE
prevalence model uses all DD-observed lots	TRUE
prevalence model keeps > 90% of DD lots	TRUE
weight/yield inconsistencies are flagged, not dropped silently	TRUE
count panels have a positive denominator	TRUE
every harvest outcome produced a model	TRUE
no correction invented a value from nothing	TRUE
originals preserved for every corrected lot	TRUE
no weight in the models is out of range	TRUE
corrected lots reproduce the recorded yield	TRUE

```
## R version 4.3.3 (2024-02-29)
## Platform: x86_64-pc-linux-gnu (64-bit)
## Running under: Ubuntu 24.04.4 LTS
##
## Matrix products: default
## BLAS:   /usr/lib/x86_64-linux-gnu/blas/libblas.so.3.12.0
## LAPACK: /usr/lib/x86_64-linux-gnu/lapack/liblapack.so.3.12.0
##
## locale:
## [1] C
##
## time zone: Etc/UTC
## tzcode source: system (glibc)
##
## attached base packages:
## [1] stats     graphics  grDevices utils     datasets  methods   base
##
## other attached packages:
##  [1] knitr_1.45          broom.mixed_0.2.9.4 emmeans_1.10.0      lme4_1.1-35.1       Matrix_1.6-5
##  [6] ggplot2_3.4.4       forcats_1.0.0       stringr_1.5.1       purrr_1.0.2         tibble_3.2.1
## [11] tidyr_1.3.1         dplyr_1.1.4
##
## loaded via a namespace (and not attached):
##  [1] gtable_0.3.4        TMB_1.9.10          xfun_0.41           bslib_0.6.1         lattice_0.22-5
##  [6] numDeriv_2016.8-1.1 vctrs_0.6.5         tools_4.3.3         generics_0.1.3      pbkrtest_0.5.2
## [11] parallel_4.3.3      sandwich_3.1-0      fansi_1.0.5         highr_0.10          pkgconfig_2.0.3
## [16] lifecycle_1.0.4     compiler_4.3.3      farver_2.1.1        munsell_0.5.0       codetools_0.2-19
## [21] htmltools_0.5.7     sass_0.4.8          yaml_2.3.8          pillar_1.9.0        furrr_0.3.1
## [26] nloptr_2.0.3        jquerylib_0.1.4     ellipsis_0.3.2      MASS_7.3-60.0.1     cachem_1.0.8
## [31] boot_1.3-30         multcomp_1.4-25     nlme_3.1-164        parallelly_1.37.1   tidyselect_1.2.0
## [36] digest_0.6.34       mvtnorm_1.2-4       stringi_1.8.3       future_1.33.1       listenv_0.9.1
## [41] labeling_0.4.3      splines_4.3.3       fastmap_1.1.1       grid_4.3.3          colorspace_2.1-0
## [46] cli_3.6.2           magrittr_2.0.3      survival_3.5-8      utf8_1.2.4          broom_1.0.5
## [51] TH.data_1.1-2       withr_2.5.0         scales_1.3.0        backports_1.4.1     estimability_1.4.1
## [56] rmarkdown_2.25      globals_0.16.2      zoo_1.8-12          coda_0.19-4.1       evaluate_0.23
## [61] glmmTMB_1.1.8       mgcv_1.9-1          rlang_1.1.3         Rcpp_1.0.12         xtable_1.8-4
## [66] glue_1.7.0          minqa_1.2.6         jsonlite_1.8.8      R6_2.5.1            multcompView_0.1-9
```