UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

The Farmer and the "Good Cow": Implications of Emerging Technology for Knowledge, Performance, and Recognition in Expert Dairy Work

Permalink

https://escholarship.org/uc/item/9ph5z3nq

ISBN

9798186386063

Author

Shen, Roni

Publication Date

2026-06-09

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

UNIVERSITY OF CALIFORNIA

Santa Barbara

The Farmer and the "Good Cow": Implications of Emerging Technology for Knowledge, Performance, and Recognition in Expert Dairy Work

A dissertation submitted in partial satisfaction of the requirements for the degree

Doctor of Philosophy in Technology Management

by

Roni Shen

Committee in charge:

Professor Paul Leonardi, Chair

Professor Matt Beane

Professor Nelson Phillips

June 2026

The dissertation of Roni Shen is approved.

Matt Beane

Nelson Phillips

Paul Leonardi, Committee Chair

May 2026

The Farmer and the “Good Cow”: Implications of Emerging Technology for Knowledge, Performance, and Recognition in Expert Dairy Work

Copyright © 2026

by

Roni Shen

For my mother,

who taught me that education is a privilege worth chasing,

and that knowledge need not come with arrogance.

ACKNOWLEDGEMENTS

First and foremost, I want to express my deepest gratitude to the farmers, herdsmen, feeders, milkers, veterinarians, technicians, and dairy owners – many of whom are first- or second-generation immigrants to the U.S. – who participated in this study. Their work is highly skilled, unglamorous, essential, and largely invisible to the public, yet they opened their (barn) doors and graciously shared the intricate realities of their work with patience and generosity to this grateful greenhorn. And, of course, thanks to the cows for providing the greatest gift to humankind: cheese.

Countless people have helped me reach this milestone, but the sad reality is that I can only list a few. Particularly, I am grateful to my advisor, Paul Leonardi, for his patience and tutelage over the last four years – for pushing me when I needed to be pushed, and for yielding when I needed time to reconstruct myself. Thanks to my committee members, Matt Beane and Nelson Phillips, for their sustained intellectual and emotional support. Thanks also to the wonderful staff of Technology Management who facilitated my academic development. Without them, this project would not have been possible. I have also been blessed with many wonderful collaborators and role models, the most influential being Angelos Kostis whose friendship and mentorship have shaped who I am as a scholar and researcher. I am grateful to him and others with whom I've had the privilege of working on this project, particularly Pam Hinds and Diane Bailey for supporting me in the production of this dissertation, as well as to Zach Tan for being ready and willing to trudge through the mud with me on the farm.

I see my PhD journey starting long before I opened that acceptance email – with my family, who gave me the solid foundation I needed to pursue such an endeavor. Thanks to my mom, Peace Song, for the impossible sacrifices she made to bring our family to the U.S. and

to give us the opportunities she never had. Thanks to my dad, Lloyd Shen, for teaching me that it's never too late to reinvent yourself. Thanks to my brother, Jun Shen, who taught me to be resilient and to charge forward with courage and integrity through all life's challenges; to my wonderful sister-in-law, Mina Kay Shen, who devotes her boundless, infinite kindness and warmth to our family every day; and to my nieces, Aria and Mila, who never let me forget to play. In a way, this dissertation reflects all the values they imbued in me – to persist through hardship without losing sight of the humanness of others, and to dignify all people regardless of their background.

During my time in the PhD program, there were few people who have had as much impact on the quality of my life as my fellow PhD students. The solidarity, the laughs, the scalding hot tea spilt over ungodly amounts of wine and cheese all enriched what might otherwise have been an isolating slog. In the order that they came into my life: thanks to Virginia Leavell for being a badass friend and role model; to Danielle Bovenberg for her infinite generosity of heart, mind, and spirit; to Dylan Kreis for being my virtual buddy through online coursework during a global pandemic; to Sienna H. Parker, my academic sibling, for the hours of commiseration/group therapy during "work" sessions at every coffee shop in town; to Sai Kalvapalle, our honorary TM student, for always being down for an impromptu Garbage concert or late night stroll; to Fares Ahmad for making the world a more compassionate place and for being an open-hearted dance partner; to Brandon Lepine for elevating the vibe of our motley crew above and below deck; to Mostafa Khoshbash for always being open to trying new things, including tarot card readings; to Danielle Morales for being unapologetically entertaining; to Nisa Gürbüz for being the best part-time housemate and thrifting companion; to Masoom Suchdeo for sharing his encyclopedic knowledge as unrestricted as his

benevolence; to Yuhan Liu and Seongho Hwang, for bringing their unique perspectives and joyful spirits to our tiny cohort, even though our paths crossed only briefly. Belonging with this group has truly been the highlight of my graduate experience.

Beyond the university, I'm grateful to my beloved Santa Barbara community that raised me since 2007, facilitating my growth from a wee undergrad to where I am today. To Kyle Peterson for supporting my moonshot dream of becoming a backyard farmer and to Jack (AKA "Jackaroo") for being the best pup and amateur chicken herder a gal could wish for. To Jonathan Wang for being my surrogate brother and bouldering buddy. To Judy Trejos for taking me in, no questions asked, when I had nowhere safe to lay my weary head. To Marly Goldberg, with whom I shared many discussions over the nature of life, love, and veterinary drama over way too much sushi. To my beloved quartet, This Is Us, consisting of Carol Dirner (our fearless lead), Diane Rothberg (our tenacious tenor), and Elaine Watson (our pint-sized bass), for inviting this vagabond baritone to create sweet harmonies together. To Mike Holt (our biggest fan and barbershop husband) who, with Elaine, are two of the best humans and tortoise lovers on the planet. I am eternally grateful to Brett Aho, who never lost sight of the light within me, even when I believed it had been snuffed out – and for reminding me that life is meant to be enjoyed. Thanks to the mental health professionals who taught me the skills and gave me the confidence to break free of old patterns that no longer served me. And of course, I'm grateful for the unlimited pets, purrs, and snuggles of my chatty feline roommate, Cool Beans. These dear friends carried me through hardships I did not know how to navigate on my own. I am truly grateful to these and many more that cannot fit into a few short pages.

Among those people are the ones who shared my first love: music. Having an outlet for embodied creativity and expression has kept me sane through more than just this PhD. I've

been lucky enough to belong to some wonderful groups that showed me the stabilizing power of art and community. Thanks to the Barbershop divas of Carpe Diem and the choral storytellers of Adelfos Ensemble (with special thanks to my fearless maestro and dear friend, Temmo Korisheli) who provided containers of camaraderie and co-creation where I could practice my musical meditation. I'm also forever grateful to the diverse community of music lovers and patrons of the arts that allow amateurs like me to perform with the world-class musicians of the Santa Barbara Symphony. Music has always been, and remains, the greatest enrichment of my life.

While music offered transcendence, the people I met along my slightly scattered career path kept me grounded. Thanks to the innumerable coworkers, clients, and counterparts I've crossed paths with – former strangers with whom I've laughed, cried, sweat, and bled. They include squid fishermen, truck drivers, dock workers, urchin divers, marine biologists, sea lion rescuers, ranchers, vaqueros, horse vets, dog trainers, winemakers, pig farmers, and all manner of humans from every walk of life. And to my veterinary medicine community: I will continue championing mental health and the imperative of human dignity for a profession that gives so much of themselves to their communities while asking so little in return. In part, this dissertation is dedicated to them.

Most importantly, thanks to my best friend and loving cheerleader, Livius Muff, who always sees the best in me, despite plenty of evidence to the contrary. I look forward to our next adventure in Switzerland… and the Swiss cows just happen to be a bonus.

This research would not be possible without the support of the National Science Foundation under Grants Project #21-616. My hope is that our collective sanity as a society will be restored soon so future aspiring researchers can benefit as I have from publicly funded science in service of the greater good of humanity.

CURRICULUM VITAE

RONI SHEN
Ph.D. Candidate in Technology Management
University of California, Santa Barbara's College of Engineering

EDUCATION

Ph.D. University of California, Santa Barbara | Technology Management (June 2026)

Committee: Paul Leonardi (Chair), Nelson Phillips, Matt Beane

B.S. University of California, Santa Barbara | Ecology, Evolution, & Marine Biology (2011)

PROFESSIONAL EMPLOYMENT

2026 Postdoctoral Researcher | ETH Zürich Department of Management, Technology, and Economics

PUBLICATIONS

Published Journal Articles, Book Chapters, & Conference Proceedings:

Leonardi, P. M., Parker, S. H., & Shen, R. (2024). How Remote Work Changes the World of Work. Annual Review of Organizational Psychology and Organizational Behavior, 11, 193-219.

Shen, R., Phillips, N. (2026). Revisiting Discourse and Institutions: Exploring the Role of Digital Technology in Institutionalization. In Research in the Sociology of Organizations: Beyond Microfoundations and Macrofoundations: A Cross-level Linguistic Perspective of Institutions. Emerald Publishing Limited.

Shen, R. (2024). Constructing Automation: Rethinking the Automation-Augmentation Paradigm in the Age of AI. Academy of Management Best Paper Proceedings.

Articles Under Review:

Kostis, A., Hinds, P., Leonardi, P., Shen, R. When Technology Surprises: Data Objects, Discrepant Events, and The Process of Serendipitous Learning. Preparing for submission to Strategic Management Journal special issue on "Chance, Luck, and Serendipity in Strategy and Management."

Manuscripts in Preparation:

Shen, R., Leonardi, P. Performing Under Threat: How Threat Intensity and Specificity Shape Experts' Use of New Technology.

Parker, S. H., Shen, R., & Leonardi, P. Burned Out by Technology: How Perceived Autonomy Shapes Strategies to Reduce Digital Exhaustion.

Shen, R., The Farmer and the Good Cow: Digital Representing as a Process of Evaluative Restructuring.

Shen, R., Expertise Inference Under Digital Visibility.

Kostis, A., Giordano, J., Shen, R., Developing AI for Real Practice: A Cross-Disciplinary Approach to Improving Predictive Models in Dairy Farming.

Shen, R., Arrested Development: Competing Logics and the Impacts of Corporate Acquisition on Expertise in Veterinary Medicine.

Shen, R., Weaponizing Practice: How Private Equity Firms Use Technology to Dismantle Professional Jurisdiction in Veterinary Medicine.

CONFERENCE PRESENTATIONS & ROUNDTABLES

Kostis, A., Hinds, P., Leonardi, P., Shen, R. (2025). When Technology Surprises: Data Objects, Discrepant Events, and The Process of Serendipitous Learning. Accepted to the Chance, Luck, and Serendipity in Management and Strategy PDW at the Academy of Management (AOM) Annual Meeting in Copenhagen, Denmark.

Shen, R. (2024). Constructing Automation: Rethinking the Automation-Augmentation Paradigm in the Age of AI. Presented at the Academy of Management (AOM) Annual Meeting in Chicago, IL, USA.

Shen, R., Kostis, A., Hinds, P., Leonardi, P. (2024). Judgment Amidst and Through Data Provocation: Producing Temporal and Relational Reflexivity. Accepted to the Process Organization Studies (PROS) Symposium in Paphos, Cyprus.

Shen, R., Leonardi, P. (2024). Algorithmic orientations: A constructivist approach to AI in organizational decision-making. Presented at the 40th European Group for Organizational Studies (EGOS) Colloquium in Milan, Italy.

Parker, S. H., Shen, R., & Leonardi, P. (2023). Burned Out by Technology: How Perceived Autonomy Shapes Strategies to Reduce Digital Exhaustion. Presented at the Academy of Management (AOM) Annual Meeting in Boston, MA, USA.

Shen, R., Parker, S. H., & Leonardi, P. (2022). Tired from Tech: Knowledge Workers' Experiences Coping with Digital Exhaustion. Presented at the International Conference on Information Systems (ICIS) in Copenhagen, Denmark.

Shen, R. (2022). The Impact of Corporate Acquisition on Professional Identity: Competing Institutional Logics in Veterinary Medicine. Presented at the People and Organizations Conference at Wharton in Philadelphia, PA, USA.

WORKSHOPS & CONFERENCES

2020-2025	Academy of Management
2025	Organizing the Climate Transition PhD Course (virtual)
2024	European Group for Organizational Studies (Milan, Italy)
2024	KIN Summer School (VU Amsterdam)
2024	Qualitative Methods Workshop (NOVA, Lisbon, Portugal)
2024	Process Organization Studies (Paphos, Cyprus)
2022	International Conference on Information Systems (Copenhagen, Denmark)
2022	Social – Symbolic Work Seminar (virtual)
2022	People and Organizations (Wharton School, Philadelphia, USA)
2022	Carnegie School of Organizational Learning Academy (virtual)
2022	Warwick Summer School on Practice and Process Studies (virtual)
2021	AOM New Doctoral Student Consortium (virtual)

RESEARCH EXPERIENCE

Management:

2023-2025 Graduate Student Researcher | Big Data on the Dairy Farm: Relational Transformations across Agricultural Occupations and Organizations with the Rise of Digital Technologies

NSF Funding: Project #21-616

PI: Paul Leonardi (UCSB), Pam Hinds (Stanford U.), Diane Bailey (Cornell U., emeritus)

2021-2022 Graduate Student Researcher | Psychological Safety and the Role of Digital Technologies in Warehousing

PI: Matt Beane (UCSB)

Life Science:

2007-2011 Research Assistant | UC Santa Barbara Marine Science Institute

Assisted in data collection, field sampling, and specimen processing for two long-term ecological monitoring programs and one Ph.D. project

PI: Steve Gaines (UCSB), Jenn Caselle (UCSB), Sarah Teck (UCSB)

TEACHING EXPERIENCE

Management:

2026 Teaching Assistant | TMP 402: Digital Marketing (1 quarter)

2022-2025 Trainer | Teaching Assistant Orientation for Technology Management

2021-2022 Teaching Assistant | TMP 120: Business Strategy (3 quarters)

Average teaching evaluation score: 1.1 (on scale of 1-5, 1 being highest)

Life Science:

2011 Instructor, Translator, & Wilderness Guide | Wildlands Studies Program, CSU Monterey Bay

AWARDS, HONORS, & GRANTS

2020-2021 Regents Fellowship ($29,000) | University of California (nominated)

2024 Academy of Management Meeting in Chicago, IL | Best Paper Designation (CTO Division)

2024 Graduate Student Association Conference Travel Grant ($500) | UC Santa Barbara

2024 UC Senate Doctoral Student Conference Travel Grant ($1,500) | University of California

SERVICE

To the Academic Community:

2025 Ad Hoc Reviewer | Academy of Management Meeting (CTO Division)

2025 Organizer | Paper Development Workshop at the Academy of Management Annual Meeting

PDW Title: “Navigating Challenges in Adopting a Relational Perspective on Data and Emerging Digital Technology”

2024-2026 Member | PhD AI & Innovation Interest Group (research consortium)

2023-2025 Supervisor | NSF Dairy Project (supervised undergraduate research assistants during field work)

2020-Present Member | Academy of Management OMT, CTO, & MOC divisions

2020-Present Participant | Technology and Organization Research Seminar (UCSB)

2021 Organizer | Psychology of Technology Institute Annual Conference

To the University Community:

2024-2026 Mentor | UCSB Graduate Scholars Program (mentor first-year PhD students in STEM)

2022-2025 Trainer | Teaching assistant orientation (facilitated new TA training for the department)

2022-Present Committee Chair & Organizer | Dean's Research Lecture Series (organized invited faculty visits)

To the Santa Barbara Community:

2019-2022 Board Member | Adelfos Ensemble (non-profit performing arts organization)

2018-2019 Ambassador | The Chamber of Commerce of the Santa Barbara Region (non-profit)

2017-2018 Manager Member | Santa Barbara Veterinary Practice Managers Group (local industry association)

INDUSTRY EXPERIENCE

2018-2020 Technician, Lead | Advanced Veterinary Specialists

2012-2018 Technician, Practice Manager, Director of Operations | Vet's Here! LLC

2012-2014 Fisheries Technician | California Department of Fish and Wildlife

2011 Manatee Photo-ID Technician | Florida Fish and Wildlife Conservation Commission

2009-2013 Rescue & Rehabilitation Supervisor | Marine Mammal Center

ABSTRACT

The Farmer and the "Good Cow": Implications of Emerging Technology for Knowledge, Performance, and Recognition in Expert Dairy Work

by Roni Shen

Across organizations and occupations, emerging technologies are reshaping expert work in ways that the literature has struggled to make sense of holistically. Research on expertise has developed along three largely separate tracks, treating expertise as knowledge held by individuals, as performance enacted for audiences, or as recognition conferred by others, and each has produced its own account of how technology disrupts expertise based on these ontological divides. However, these are only partial accounts of a larger phenomenon that is difficult to capture in one study. This dissertation attempts to do just that. I argue that knowing, performing, and recognizing are not competing definitions of expertise but component, interwoven constitutive processes within a unified ecological framework of expert work, and that technology, once it enters expert work, becomes endogenous to each of them and to the recursive loops that bind them together. I develop this argument through three qualitative field studies of expert work in the U.S. dairy industry, where farmers are navigating the rapid integration of robotic milkers, activity-tracking collars, genomic testing, predictive algorithms, and integrated data platforms into work that was historically organized around embodied judgment and "cow sense." Drawing on participant observations of, and interviews with, seventy-four dairy workers and twelve large animal veterinarians in California across sixteen dairy farms, two calf ranches, and five large animal veterinary practices, I show how these three constitutive processes of expertise are recursive and interdependent. Study 1

examines knowing as a recursive process of representing an object of expertise in digital form. Findings show how digital systems rendered cows as four different evaluative objects – an inscription-compatible body, a biographical record, an algorithmic prediction, and a comparative position in a herd distribution – and how farmers' ongoing congruence work to align digital and physical realities cumulatively reshaped both digital and physical objects over time. Study 2 examines performing as a recursive process of orienting towards technologies in ways that are informed by non-technological interactions. Findings show how farmers' experiences of expertise threat from employers, peers, and subordinates became the interpretive lens through which the same technology was viewed either as a tool, a symbol, or a control mechanism. Study 3 examines recognizing as a recursive process of inferring expertise from digital traces of work. Findings show how digital visibility gave dairy veterinarians insights into farmers' backstage activities that were not meant to be seen. Vets arrived at inferences about gaps in farmers' knowledge, execution, or values, and calibrated interventions to fill these gaps while buffering their strategies to protect their working relationships with farmers. Read together, the three studies advance an integrated framework in which technology participates in expertise not by replacing knowledge, mediating performance, or shifting recognition criteria as discrete effects, but by entering the entangled relationships through which all three processes unfold. The dissertation contributes to scholarship on expertise by bridging long-standing ontological divisions across the field, and to scholarship on technology and organizing by specifying how technologies become endogenous to expert work in the densely relational, interdependent settings that characterize modern technologically saturated organizations.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS v

CURRICULUM VITAE x

ABSTRACT xv

LIST OF TABLES xxv

LIST OF FIGURES xxvi

CHAPTER 1: INTRODUCTION 1

Perspectives on Expertise and Technological Change 3

What is Expertise? 3

Expertise as Knowledge 5

Technology Changes What Experts Know 6

Expertise as Performance 8

Technology Changes How Experts Perform 9

Expertise as Recognition 11

Technology Changes How and When Experts are Recognized 12

Research Questions 13

Overview of Three Empirical Studies 14

Study 1: Knowing as a Recursive Process of Representing – How Data-Intensive Technologies Reconstitute the "Good Cow" 15

Key Findings and Contributions 15

Study 2: Performing as a Recursive Process of Orienting – How Pre-Existing Performance Threats Shape Technology Use in Practice 16
Key Findings and Contributions 16
Study 3: Recognizing as a Recursive Process of Inferring – How Audiences Construct Expertise Inferences Under Digital Visibility 17
Key Findings and Contributions 18
CHAPTER 2: METHODS 19
Studying Expertise and Technology in Organizations 19
A (Far Too) Brief Overview of Dairy Farming 20
Activities Central to Dairy Operations 23
Expert Roles on the Farm 25
The Technological Evolution of the Dairy Industry 27
The Modern Dairy Farm 33
Dairy Farming as an Ideal Research Context 37
Research Design and Epistemological Approach 40
Site Selection 41
Data Collection 42
Data Analysis 45
Study 1: Knowing as a Recursive Process of Representing – How Data-Intensive Technologies Reconstitute the "Good Cow" 46

Study 2: Performing as a Recursive Process of Orienting – How Pre-Existing Performance Threats Shape Technology Use in Practice 47

Study 3: Recognizing as a Recursive Process of Inferring – How Audiences Construct Expertise Inferences Under Digital Visibility 48

CHAPTER 3: Knowing as a Recursive Process of Representing – How Data-Intensive Technologies Reconstitute the "Good Cow" 50

Digital Representations and the Physical World 52

Working *With* Representations 53

Working *Towards* Representation 56

From *Representation* to *Representing* 59

Methods 62

Research Setting 62

Site Selection and Technologies 64

Data Collection 65

Data Analysis 66

Findings 68

Inscription Bottleneck: Representing requires cow-robot compatibility 70

Representability Restructures Evaluations of a "Good Cow" 73

Congruence Work Aligns Bodies and Robots 75

Consequences of the Inscription Bottleneck 77

Biographical Rendering: Representing expands what is visible....78
Clean, Comprehensive Records Restructure Evaluations of a "Good Cow"....80
Congruence Work Aligns Embodied Observations and Digital Insights....82
Consequences of Biographical Rendering....83
Algorithmic Rendering: Representing Imbues Data with Meaning....84
Promising Futures Restructure Evaluations of a "Good Cow"....86
Congruence Work Aligns Present Realities and Predicted Futures....87
Consequences of Algorithmic Rendering....90
Comparative Rendering: Representing Makes Individuals Commensurable....91
Comparison Restructures Evaluations of a "Good Cow"....92
Congruence Work Aligns Herd Composition and Representational Goals....94
Consequences of Comparative Rendering....95
Summarizing the Representing Process....97
Discussion....100
Representing as Ongoing Process....101
Evaluative Restructuring and the Ongoing Redefinition of "Goodness"....103
Congruence Work and the Materialization of the Epistemic....107
Conclusion....111
CHAPTER 4: Performing as a Recursive Process of Orienting – How Pre-Existing Performance Threats Shape Technology Use in Practice....113

Perspectives on Technological Change in Expert Work 115
Technology as a Threat to Expert Stability 122
Technology as the Starting Point for Interpretation 126
Methods 128
Research Setting and Site Selection 128
Data Collection 130
Data Analysis 133
Findings 138
DF6: Responding to Low Specificity, Low Intensity Threat 139
Consequences of a Critical Orientation for Expertise 145
DF11: Responding to Low Specificity, High Intensity Threat 146
Consequences of a Deliberative Orientation for Expertise 150
DF9: Responding to High Specificity, Low Intensity Threat 151
Consequences of a Deferential Orientation for Expertise 155
DF2: Responding to High Specificity, High Intensity Threat 157
Consequences of an Appropriative Orientation for Expertise 161
Discussion 163
Interpreting Technology through a Lens of Threat 164
Rethinking the Link Between Technology Use and Expertise Performances 168
Conclusion 171

CHAPTER 5: Recognizing as a Recursive Process of Inferring – How Audiences Construct Expertise Inferences Under Digital Visibility 173

Expert Work Under Digital Visibility 176

Expertise as Performed and Socially Attributed 176

Opacity Enables Expert Image Regulation 178

Digital Visibility into the Backstage 180

Methods 184

Research Setting 184

Data Collection 187

Data Analysis 189

Findings 191

Glimpsing the Backstage 191

Component 1: Signal Recognition within Existing Interpretive Frames 194

How Certain Signals Become Salient 197

Component 2: Disambiguation through Two Channels 198

Channel 1: Relational Contextualization 199

Channel 2: Data Interrogation 200

Disambiguation Under Constraint 203

Component 3: Inference as Tentative Commitment 204

Knowledge Inferences 205

Execution Inferences....207
Values Inferences....209
Component 4: Calibrated Intervention....211
Intervening to Fill Knowledge Gaps....213
Intervening to Fill Execution Gaps....215
Intervening to Fill Values Gaps....217
The Recursive Inference Cycle....219
Discussion....221
A Theory of Inference Under Digital Visibility....223
The Changing Conditions of Expert Attribution....228
Inference as Recursive Work....230
Conclusion....233
CHAPTER 6: CONCLUSION....235
Study 1: Knowing as Embedded in Performance and Recognition....236
Study 2: Performing as Translating Knowing to Recognition....238
Study 3: Recognition as Both Reading and Conditioning of Expert Work....240
Integrating the Three Studies....242
The Limitation of Single-Process Accounts....244
Expertise as a Relational Ecology....246
Implications....248

For the Study of Expertise 248

For the Study of Technology and Organizing 249

For Practice 250

Limitations and Future Directions 251

A Closing Note on the Worker 253

REFERENCES 257

APPENDIX A: Conventional versus Data-intensive Dairy Operations 277

APPENDIX B: Table of Field Sites 281

APPENDIX C: Table of Informants 283

APPENDIX D: Images from the Field 287

Images of Work 287

Images of Milking Systems 292

Images of Data Systems and Platforms 294

Images of Cows and the People Who Love them 299

LIST OF TABLES

Table 1.1. Three dominant conceptualizations of expertise in management literature 4

Table 2.1. Summary of collected data 44

Table 3.1. Overview of evaluative restructuring of a "good cow" 96

Table 4.1. Summary of literature on technological threats and expert response 116

Table 4.2. Groups of mechanisms through which technology disrupts expertise 120

Table 4.3. Characterization of threat and how the character of threat shapes technology interpretation and use 136

Table 5.1. Summary of informant veterinarians 188

LIST OF FIGURES

Figure 2.1. Image of an RFID ear tag on a Holstein calf 29

Figure 2.2. Image of a rotary milking parlor. 30

Figure 2.3. Image of a Holstein cow wearing an activity tracking collar 31

Figure 2.4. Image of a cow being milked in a Lely A5 Astronaut robotic milker 32

Figure 2.5. Screenshot from the Zoetis Genomics website showing the marketed use case for genomics testing for breeding decision support..... 33

Figure 2.6. An example of a data map from DF11 depicting all the data-generating devices and the linkages between human, technological, and physical entities through which data flows..... 35

Figure 3.1. Model of the representing process 70

Figure 3.2. Moments in the representing process that invoked evaluative restructuring and congruence work 98

Figure 3.3. Model of ontological co-evolution over time facilitated by congruence work 99

Figure 3.4. Evaluative restructuring and congruence work via evaluative moments in the representing process 101

Figure 4.1. Studies of expertise in relation to the broader temporal context of expertise performance 123

Figure 4.2. Relational diagrams at four farms..... 135

Figure 4.3. Descriptive model of threat experiences, technological orientations, threat mitigation practices, and outcomes for expertise.....139

Figure 4.4. Theoretical model of compensatory technology adoption in response to expertise threat 164

Figure 5.1. A model of recursive expertise inference.....194

Figure 5.2. Model of expertise inference under digital visibility.....222

Figure 6.1. Integrated model of the three interrelated recursive processes that constitute expertise, and how technologies become endogenous to each 244

CHAPTER 1: INTRODUCTION

Across occupations and industries, the introduction of new technologies has raised pressing questions about the future of expert work. As technological capabilities advance at a staggering pace, scholars worry that they are encroaching increasingly into domains that have previously been the exclusive purview of human experts (Brynjolfsson & Mitchell, 2017; Faraj et al., 2018; Frey & Osborne, 2017; von Krogh, 2018). This is not a new phenomenon, however. Scholars of work and organizations have long been interested in how technologies alter what and how experts know (Alaimo & Kallinikos, 2021; Østerlie & Monteiro, 2020; Pachidi et al., 2021), what they do (Bailey et al., 2012; Beane, 2019), and how they engage with one another (Barley, 1986, 1990; Beane & Anthony, 2023).

Despite decades of theorizing and empirical findings that show the complexities of expertise, the field continues to be organized around a set of foundational debates about what expertise *is* and whether it can be displaced, augmented, or rendered obsolete by technology (Acemoglu & Restrepo, 2019; Brynjolfsson & McAfee, 2014; Lei & Kim, 2024; Pakarinen & Huising, 2023). These debates persist in part because the expertise literature has developed along three distinct analytic tracks. One track treats expertise as substantive knowledge or experience-based skill held by individuals and asks whether technology can replicate or replace those capabilities. A second treats expertise as situated performances enacted through social interface and asks whether technology disrupts the relational conditions under which these interactions take place. A third treats expertise as a social attribution conferred by audiences and asks whether technology reshapes who is granted the status of expert and how expertise is recognized. Each of these traditions has produced rich and diverse contributions to our understanding of what constitutes expertise in our society, but the cumulative effect of

ontological divergence has been aggregation around three dominant core assumptions about how technology interacts with expertise and expert work.

This dissertation shows that these three processes (knowing, performing, and recognizing) are not separate accounts of expertise to be chosen among, but coexisting and mutually constitutive phases of a single ongoing achievement, with technology as an endogenous participant in each. Bringing the three dominant treatments together makes visible what each, working alone, has tended to obscure: the recursive relations through which a change in one process moves through the others, and the role technology plays once it becomes embedded in those relations. The framing shifts from technology as an exogenous disruptor of expert work to technology as an entangled component of the recursive processes through which expert work is continuously constituted (Bailey et al., 2022). Expertise in modern organizations is embedded in relational networks of interdependent expert work – networks in which what experts know, how they perform that knowing for others, and how others come to recognize them as experts are not separable phenomena but mutually constitutive ones. When technology enters such networks, it becomes part of how experts know, how they perform, and how they are recognized – and, crucially, part of the recursive loops that bind these processes together.

Drawing on three empirical studies of expert work in data-intensive dairy farming, the dissertation shows that *knowing* is recursive in that digital representations continuously redefine and reshape the physical referents they depict; that *performing* is recursive in that expert performances are shaped by ongoing engagements with audiences which technologies then enter and reshape; and that *recognizing* is recursive in that audience inferences are conditioned by perceptions of expertise based on expert work rendered visible through digital tools. Read together, these three recursive processes feed into one another. Technology's role

is not to disrupt stable expertise from outside but to participate in, and reshape, the embedded relational dynamics through which expertise is continuously produced.

This integrated framing has consequences both for how we think about expertise and how we think about technology. For expertise scholars, it offers a way to move beyond ontological disputes about what expertise *is* and toward empirical questions about how the processes that constitute expertise are linked. For scholars of technology and organizing, it offers a more granular account of how technologies become endogenous to expert work, not by replacing knowledge, mediating performance, or shifting recognition criteria as discrete effects, but by entering the recursive relationships through which all three of these processes unfold.

Perspectives on Expertise and Technological Change

What is Expertise?

To explain how technology becomes a co-constituent of expertise, it is important to first understand how the expertise literature describes what it is. However, what expertise is depends on who you ask (Heimstädt et al., 2024). Psychologists and management scholars who tend towards a more cognitive approach to expertise claim that expertise is a substantive ability that resides in the individual and is accrued through deliberate practice (Ericsson et al., 1993), accumulation of experience (Dreyfus & Dreyfus, 2005), or structured study or training (Gobet & Simon, 1996). By comparison, those who adopt a relational ontology argue that expertise is instantiated and enacted within interactions (Eyal, 2013; Pakarinen & Huising, 2023), while those who take a more social constructivist stance look towards institutional signals and culturally legitimized indicators of expertise, such as degrees, credentials, licensure, and

professional affiliations that legitimize expertise through social recognition and conferral (Abbott, 1988).

From these distinct ontological foundations, management research on expertise has diverged into several self-referential scholarly communities. Within each, as Anteby and Holm (2021, p. 312) observed, "the terms expert and expertise are generally considered self-explanatory" to those who share similar core assumptions of what constitutes expertise. This is problematic because, although these communities share an interest in the lexical concept of expertise, they take it to mean entirely different things, and there has been surprisingly little exchange between them as a result. **Table 1.1** summarizes three dominant ontologies that have emerged from this fragmentation. It is offered not as an exhaustive review but as an illustration of the analytic boundaries that structure most expertise research. Each ontological approach treats expertise as a different kind of 'thing' (e.g., an attribute, an enactment, or a status) and each generates a corresponding view of technology's role. Read in isolation, each perspective produces its own kind of debate about whether expertise can survive technological change. The technology and work literature has tended to mirror this fragmentation, documenting technology's impact along these same lines. The argument advanced here is that these are not competing accounts of what expertise is. They are partial views of three interrelated processes that together constitute expertise in practice.

Table 1.1. Three dominant conceptualizations of expertise in management literature.

Perspective on Expertise	Definition under this perspective	Characteristics of expertise	Intellectual foundations	Technology's role in expert work
Knowledge	"Elite, peak, or exceptionally high levels of performance	Individual level Substantive ability, efficacy, intuition,	Cognitive science and psychology Substantive ontology	Technology can codify human expertise and changes what/how experts know

	on a particular task or within a given domain." (Bourne et al., 2014, p. 1)	judgment; accrued knowledge, or skill Expertise is possessed, stable and measurable	(Dreyfus & Dreyfus, 2005; Ericsson et al., 1993; Gobet & Simon, 1996)	(Alaimo & Kallinikos, 2021; Beane, 2019; Lebovitz et al., 2022)
Performance	"A network linking together agents, devices, concepts, and institutional and spatial arrangements" (Eyal, 2013, p. 863)	Interactional level Situated interactions, communications, practices, or performances Expertise is discursive, emergent, and ephemeral	Sociology of expertise Relational ontology (DiBenigno, 2020; Eyal, 2013; Huising, 2015; Treem, 2012)	Technology can disrupt relationships and changes what/how experts do (Anthony, 2021; Bailey et al., 2012; Beane & Anthony, 2023; Leonardi & Treem, 2012; Treem, 2013)
Recognition	"Professions are exclusive occupational groups applying somewhat abstract knowledge to particular cases" (Abbott, 1988, p. 8, speaking of expertise in professions)	Social level Socially constructed, conferred, and recognized by others Expertise is attributional, contingent, and negotiated	Sociology of profession Social constructivist ontology (Abbott, 1988; Bourgoin & Harvey, 2018; Bunderson, 2003; Su, 2012; Treem & Leonardi, 2017)	Technology changes the basis upon which expertise is perceived, legitimized, and valued (Bechky, 2020; Botelho, 2024; Leonardi, 2007)

Expertise as Knowledge

A knowledge-based perspective rooted in cognitive science and psychology conceptualizes expertise as substantive knowledge, intellective ability, and cognitive schemas held within the individual. Knowledge, in this sense, serves as the foundational resource of expertise and can be gained, accumulated, and developed through deliberate practice and experience (e.g., Dreyfus & Dreyfus, 2005; Ericsson et al., 1993; Gobet & Simon, 1996). Sometimes called a realist (Heimstädt et al., 2024) or substantialist perspective (Pakarinen &

Huising, 2023), this approach treats expertise as something that experts possess – a repository of domain-specific knowledge, problem-solving heuristics, experientially-tuned intuition, and pattern recognition capabilities that distinguish experts from novices (Chase & Simon, 1973). From this perspective, expertise is an individual attribute that increases as people accumulate experience and refine their mental models of their domain. Expert knowledge includes both explicit, codifiable knowledge and tacit knowledge and intuition that is difficult to articulate but evident in practice (Hadjimicheal et al., 2024; Nonaka & Von Krogh, 2009). Expertise, under this view, is predominantly seen as intellectual work and emphasizes what experts know about their domain, how they organize that knowledge in memory, and how they deploy it to solve problems and make decisions as the foundation of expertise.

While this perspective on expertise has shaped thinking around organizing around knowledge (Simon, 1972), recent organizational scholarship has begun to push against this individualist framing by attending to how knowledge is enacted through ongoing practice rather than possessed as a stable attribute (Orlikowski, 2002). This processual reframing of knowing in practice – from a thing one has to a thing one does – opens the possibility that technology participates in knowing not by adding to or subtracting from an individual's knowledge stock but by becoming entangled in the practices through which knowing happens.

Technology Changes What Experts Know

The majority of studies of expertise in organizations adopt a knowledge-centered view that locates expertise within individual members as "know-what" or "know-how" (Rousseau & Stouten, 2025). As such, most studies of technological change examine how technology alters what experts know and how they come to know it. Technologies introduce new ways of sensing, classifying, and interpreting phenomena in expert domains, fundamentally reshaping

the epistemic foundations of expertise (Alaimo & Kallinikos, 2022). Some technologies redistribute knowledge across organizational members, making previously specialized knowledge more widely accessible and thereby undermining the exclusivity that sustained expert status (Leonardi, 2007). Others have shown how new technologies become barriers to traditional pathways of developing expertise. For example, Beane (2019) shows how surgical robots negated traditional methods of developing surgical skills, requiring trainees to learn entirely new skills to remain effective. Other technologies redefine the cognitive foundations of expertise by digitizing tacit knowledge and transforming situated understanding into standardized inputs, often stripping away the contextual nuances that experts rely upon for judgment. Brayne's (2017) study of predictive policing, Zuboff's (1988) analysis of automated paper manufacturing, and Alaimo and Kallinikos's (2021) examination of algorithmic management all demonstrate how technologies can formalize and codify knowledge that was previously embodied in expert practice. Similarly, Lebovitz and colleagues (2021) show how diagnostic AI models in hospitals codify expert "know-what" (i.e., the declarative knowledge of medical facts) without incorporating the contextualized "know-how" that physicians develop through clinical experience. Across these studies, technology emerges as a force that can enhance, undermine, or reconfigure the knowledge foundations that underlie expertise, fundamentally altering what it means to know something as an expert.

A particularly important line of work, drawing on the science and technology studies (STS) tradition, demonstrates that technologies do not merely add new information to what experts already know; they participate in representing the objects of expertise as codified, quantified simulacra – and in so doing, functionally constitute what those objects come to "be." Latour and Woolgar (1986) showed how laboratory instruments do not simply record but

actively produce the phenomena they measure. Mol (2000) argued that the practices of measuring blood glucose and rendering a diagnosis enact diabetes by influencing what patients eat and how they perceive themselves. Heuts and Mol (2013) made a parallel observation about what constitutes a "good tomato," showing that different registers of value emerged across different practices, and that informants did not merely evaluate tomatoes against a fixed standard but actively worked to *make* tomatoes good. Goodwin (1995) also demonstrated that the tools available for quantifying physical phenomena constitute what those phenomena are as objects of inquiry. Across this tradition, representing emerges not as a neutral act of encoding pre-existing reality but as a constitutive process through which the represented object is brought into being.

Bringing this more sociological lineage into conversation with organizational scholarship on data-intensive work (Bowker & Star, 1999; Christin, 2020) reveals a recursive dynamic that neither literature has fully developed on its own. Continuous, multi-layered digital representational systems do not produce stable representations that then exert effects. They sustain ongoing processes of representing in which digital renderings and physical referents reciprocally reshape one another over time.

Expertise as Performance

Drawing from the sociology of expertise literature, a performance-based perspective conceptualizes expertise not as something that is possessed but as something that is enacted, communicated, and recognized through interaction (e.g., DiBenigno, 2020; Eyal, 2013; Huising, 2015; Treem, 2012). From this view, rooted in Goffman's (1956) dramaturgical framework showing how identity is performed for audiences in ways that manage impressions and sustain legitimacy, expertise is relationally instantiated and must be demonstrated in

specific situations to be meaningful. Under this view, simply possessing knowledge is not enough. Experts must perform that knowledge in ways that others can observe and evaluate through communicative, embodied, and material displays. This perspective emphasizes that expertise is accomplished through the strategic use of communicative resources (Chan & Hedden, 2021; Treem, 2012), material artifacts (Treem, 2013), and situated practices (DiBenigno, 2020; Huising, 2015; Orlikowski, 2000) that make expertise visible and credible to relevant audiences (Bourgoin et al., 2020; Bourgoin & Harvey, 2018). As Hughes (1958) described, experts represent *claims to competence* that are continuously evaluated, while Strauss (1978) emphasized that professional work and authority are negotiated orders maintained through everyday interactions rather than fixed hierarchies. These claims to expertise constitute how, when, and for what individuals want to be recognized for as experts (Evans, 2021; Fayard et al., 2017; Nelsen & Barley, 1997).

Technology Changes How Experts Perform

Studies of technological change that adopt a performance lens examine how technology reconfigures what experts do and how expertise is performed. Most studies in this stream focus on how technology disrupts established performance repertoires by interrupting established scripts (Barley, 1986), reconfiguring choreographed routines (Edmondson et al., 2001), and increasing the visibility of decisions that were previously inscrutable to non-experts (Karunakaran et al., 2022). For instance, Sergeeva and colleagues' (2020) study of surgical robots showed how they transformed the embodied practices through which surgical expertise was performed and required experts to adopt new collaborative practices that updated their performance repertoires. Particularly notable is Beane and Anthony's (2023) demonstration of how senior bankers and surgeons adopted "inverted apprenticeship" relationships with junior

colleagues to update their technological expertise while maintaining their performances as more experienced experts. Technology also plays a role as symbolic resource for expertise performances (Prasad, 1993). For example, Leonardi and Treem showed how information technology can act as "*stages* upon which individuals enact performances of expertise" for audiences (2012, p. 37, emphasis in original), while Pachidi and colleagues (2021) showed how sales associates symbolically adopted new algorithmic technologies in order to convince technologists that they incorporated it into their work when in fact they had not. Experts routinely use material objects and tools to enhance the persuasiveness of their claims to expertise, with technologies becoming resources that experts incorporate into their performances, shaping not just what they know but how they demonstrate that knowledge in practice (Orlikowski, 2002), how they convey authority (Bourgoin et al., 2020), and how they communicate expert status (Treem, 2016).

A performance perspective treats expertise as something continuously enacted and negotiated through interactions with audiences (Collins & Evans, 2007; Pakarinen & Huising, 2023; Treem & Leonardi, 2017). Pinch (2010) explicitly theorizes technologies as "props" in these Goffmanian performances, participating in the ongoing work of constructing and maintaining expert identities rather than threatening them from outside. This framing has important implications: if expertise is enacted through performance, then experts are constantly defending and reaffirming their standing; and the "fragility of experts" (Kang & Kim, 2022) is a routine feature of expert work rather than a consequence of technological disruption. Management consultants face credibility threats when learning unfamiliar client contexts (Bourgoin & Harvey, 2018); career advisors face legitimacy threats when client values diverge from professional guidance (Chan & Hedden, 2021); experts across fields face jurisdictional

threats from competing occupations (Lifshitz-Assaf, 2018; Truelove & Kellogg, 2016). Technologies enter these already-active performances, and they do so as resources that experts can mobilize in response to specific challenges to their standing.

Expertise as Recognition

Rooted in the sociology of professions, a recognition-based perspective emphasizes that expertise is not simply possessed or performed but must be socially recognized and conferred by relevant audiences to be functionally meaningful (Abbott, 1988; Evans, 2021). This view highlights that expertise is ultimately a social status granted through collective processes of evaluation and legitimation. Where cognitive psychology defines expertise as superior individual performance, Abbott argues that expertise is fundamentally a socially legitimated claim to jurisdiction grounded in abstract knowledge, arguing that "the thing to be explained is how expertise is structured in society, the division of expert labor" (1988, p. 359). Experts in the professions achieve authority over work domains through boundary disputes in which occupational groups compete for jurisdictional control by convincing audiences of the validity of their expertise claims, constructing jurisdictions through rhetorical, symbolic, and discursive work that draws lines between expert and lay knowledge (Gieryn, 1983). Recognition can occur through formal institutional mechanisms such as licensing, credentials, and educational requirements, or through informal processes of social attribution as audiences observe and interpret expert performances, drawing on multiple sources including professional titles, organizational affiliations, and the material trappings of expert work (Rousseau & Stouten, 2025; Vallas, 2001). This perspective underscores that audiences draw from multiple sources when forming evaluations of expertise, and that recognition is not automatic but depends on cultural understandings of what counts as expertise in a given domain.

Alongside this institutional account of recognition, a more recent line of work has examined the micro-interactional conditions under which audiences attribute expertise to specific others through observation and interpretation. Building on Goffman's (1956) distinction between front stage and backstage, scholars have shown that expert authority depends on the curated visibility of certain practices and the protected opacity of others (Anteby, 2013; Bernstein, 2012; Pratt et al., 2006; Treem, 2012). Audiences attribute expertise based on what crosses the frontstage/backstage boundary – and that boundary has historically been controlled by the expert.

Technology Changes How and When Experts are Recognized

Studies of technological change that take a recognition-based perspective explore how technology shapes the criteria by which expertise is socially evaluated and how expertise is understood and valued by different audiences. Technologies, under this view, can provide audiences with new insights into information (Leonardi, 2018), behavior (Leonardi & Treem, 2020), and others' evaluations (Orlikowski & Scott, 2014) in ways that potentially change how they perceive and evaluate experts. For example, Karunakaran and colleagues' (2022) study of social media showed how digital transparency gave audiences unprecedented visibility into expert service industries, altering the bases of evaluation while redefining the services these organizations provided. Other studies document how technological capabilities can call into question the value of traditional expertise, eroding the authority, autonomy, and discretion that experts previously enjoyed. Lifshitz-Assaf (2018) demonstrated how crowdsourcing platforms challenged NASA scientists' exclusive claims to problem-solving expertise by showing that outsiders could contribute valuable solutions. Similarly, Bechky (2020) showed how the development of DNA sequencing undermined forensic experts' authority and credibility in

court by setting a new ‘‘gold standard’’ of forensic evidence, revealing how new technologies can shift the social worth attributed to different forms of expertise, revaluing some skills while devaluing others. Across these studies, technology emerged as a participant in the cultural legitimation processes through which expertise is validated, shaping whose knowledge counts and under what conditions expertise is recognized and rewarded.

A particularly consequential structural shift involves what Leonardi and Treem (2020) call *behavioral visibility* – the digitization of work means that activities once leaving no trace now generate persistent digital records as a routine byproduct of the tools used to perform them. This shift has reorganized the conditions of expertise recognition in ways literature is only beginning to engage. The backstage of expert work is no longer fully opaque to audiences; digital traces leak across the boundary the expert once controlled. But these traces are abstracted from the reasoning, context, and tacit judgment that produced them, leaving audiences in a new position: they have visibility into the backstage without transparency into it, and the gap between what is visible and what is interpretable must be filled by inference. Recognition under digital visibility is thus a recursive process – prior inferences condition what counts as signal in subsequent encounters, audience interpretations feed back to revise prior inferences, and the working relationships between observers and observed shape what inferences become possible.

Research Questions

Taken together, these three dominant streams of literature recognize three distinct constitutions of expertise that are ontologically separated in the literature. However, empirical findings demonstrate repeatedly how knowledge, performance, and recognition are indeed distinct, but not separate. The impact of technological change, therefore, cannot be understood

from examining its interaction with one dimension in isolation. Rather, a holistic approach is necessary to reveal how technology participates in all three constitutive processes in ways the literature has yet to describe. This raises the central research question motivating this dissertation:

> *How do technologies become endogenous and co-constitutive of expertise – across the dimensions of knowing, performing, and recognizing – in modern organizations where expert work is densely relational and interdependent?*

More specifically, this dissertation examines this question in three parts:

> **Study 1:** *How do digital representations come to reshape not only what we know about the physical world, but also the physical world itself?*
>
> **Study 2:** *How do experts' ongoing performances of expertise and prior experiences of threat shape their interpretations of new technologies, and what are the implications for how they subsequently engage with these tools over time?*
>
> **Study 3:** *How do audiences infer expertise when they gain visibility into ambiguous, digitized traces of backstage work?*

I address these questions in three empirical chapters, after which I will conclude with an exploration of a fundamental, overarching question: *How do these three recursive processes, taken together, constitute and reconstitute expertise in modern expert work systems?*

Overview of Three Empirical Studies

This dissertation comprises three empirical studies, each examining one of the three interrelated recursive processes through which expertise is constituted, and each developing a framework for understanding how technology becomes endogenous to that process. Together, the three studies advance an integrated model wherein each demonstrates that a different

constitutive process is itself recursive. Read together, they show how those recursive processes link up to continuously (re)produce expertise.

Study 1: Knowing as a Recursive Process of Representing – How Data-Intensive Technologies Reconstitute the "Good Cow"

The first study develops a framework of *representing as an ongoing recursive process* through which digital technologies and the physical objects of knowing reciprocally reshape one another over time. Where existing research has tended to treat digital representations as relatively stable artifacts that, once produced, exert effects on knowing and decision-making, this study reconceptualizes representing as a continuous, distributed, recursive process – paralleling Orlikowski's (2002) shift from knowledge to knowing. The "representation" of a cow as a coherent, stable object never truly arrives; what farmers engage with is the ongoing act of representing through which both digital renderings and physical animals are continuously produced and revised.

Key Findings and Contributions

This study identifies four moments in the representing process: the *inscription bottleneck* (the cow must be made physically and dispositionally compatible with the device before any data can be produced), the *biographical rendering* (the cow as a longitudinal record), the *algorithmic rendering* (the cow as a predicted trajectory), and the *comparative rendering* (the cow as a position within a herd distribution). Each moment renders the same physical cow as a different evaluative object and invokes its own distinct criteria for what counts as a "good cow."

This study also develops the concept of congruence work – farmers' deliberate efforts to align digital and physical realities – and shows that congruence work operates differently at

different moments in the representing process. When the rendering is more plastic, farmers modify the data; when the physical referent is more plastic, farmers modify the cow. Over time, this congruence work cumulatively reshapes both digital records and physical herds.

Finally, this study proposes a model of ontological co-evolution in which representations and the physical entities they depict reciprocally transform each other, mediated by farmers' congruence work. New ways of knowing produce new physical realities, which then require new representational work, which in turn produces further physical change.

Study 2: Performing as a Recursive Process of Orienting – How Pre-Existing Performance Threats Shape Technology Use in Practice

The second study develops a framework of *technological orientations through a lens of threat* through which experts continuously navigate threats to their standing, and through which technologies become incorporated into ongoing performances as both instrumental tools and symbolic resources. The study challenges three assumptions that have shaped much of the literature on technology and expert work: a left-censored analytic frame that begins at the moment of implementation and obscures relevant prehistories (Leonardi & Barley, 2010); a stable-expertise assumption that treats pre-implementation expertise as uncontested when in fact it is continuously performed and negotiated; and a technology-first model of interpretation that locates appraisal of threat or opportunity in features encountered after implementation, when in fact appraisals begin well before the technology arrives.

Key Findings and Contributions

This study identifies pre-existing expertise threats as a primary interpretive lens through which experts engage with new technologies. Experts experiencing high-intensity, salient threats to their standing immediately enlist new technologies as resources to address

their most pressing vulnerabilities, while experts in relatively stable performance contexts orient toward the same technologies with greater skepticism or as potential threats themselves.

This study also develops a typology of threat character along two threat dimensions (intensity and specificity) that systematically explains variation in how experts interpret and incorporate technologies into their performances. The same technological features trigger very different threat/opportunity categorizations depending on experts' particular performance profiles.

Finally, this study advances a theoretical model of compensatory technology adoption, in which technology orientations function as compensatory strategies that address gaps between experts' desired professional identities and the external pressures they face.

Study 3: Recognizing as a Recursive Process of Inferring – How Audiences Construct Expertise Inferences Under Digital Visibility

The third study develops a framework of *expertise inference under digital visibility* through which audiences construct interpretations of expert competence from ambiguous digital traces of work the expert did not produce for evaluation. The study engages a structural shift in the conditions of expert recognition: the digitization of work has made the backstage of expert practice increasingly visible to audiences, but the traces produced are abstracted from the reasoning, context, and tacit judgment that generated them. Audiences gain interpretive agency while losing the interpretive structure that expert curation once supplied. The expertise literature has examined how experts perform competence for audiences but has assumed conditions of relative backstage privacy; the visibility literature has examined what digital traces reveal but has not engaged in detail with how that inferential burden gets met when expertise is what audiences are inferring.

Key Findings and Contributions

This study identifies a recursive cycle of expertise inference with four interlinked components: *signal recognition* (audiences detect digital signals carrying implicit information about the expert), *disambiguation* (audiences narrow the candidate meanings of signals through situated engagement and independent data interrogation), *inference* (audiences commit to working interpretations of the expert's competence), and *calibrated intervention* (audiences act on inferences through interventions matched to the relational risk each one carries).

This study also identifies three dimensions of inferred expertise gaps: *knowledge* (the expert lacks information), *execution* (the expert knows what to do but is not doing it reliably), and *values* (the expert's priorities diverge from what the work demands). Each carry different levels of relational risk. Audiences default to the most generous interpretation (knowledge) and revise upward toward execution and values only as more generous interpretations are ruled out.

Finally, this study develops the concept of relational buffering – a meta-practice through which audiences soften the directness of their interventions to protect the working relationships through which all subsequent inferential work flows. The form buffering takes varies systematically with the relational risk of the inference being acted on, highlighting the relational substrate of inferential work: under digital visibility, audiences hold new interpretive agency, but accurate inference depends on the working relationships that provide contextual access for disambiguation. Without that relational substrate, interpretive agency does not produce accurate inference but rather solidifies into durable misreadings.

CHAPTER 2: METHODS

Studying Expertise and Technology in Organizations

Studying how technology shapes expert work requires a context in which two conditions are simultaneously present: a domain in which expertise exists, and a set of technologies that materially alter the conditions under which that expertise is performed. Neither condition is particularly rare. Expertise, however one defines it, can be observed wherever specialized knowledge is brought to bear on a task – whether the speaker of a language, the diagnostician of a disease, or the maker of a pot (Collins & Evans, 2007). In organizations, experts are those who possess the specialized knowledge required to perform consequential work (Abbott, 1988), and every organization contains experts of some kind, varying in domain and degree. Likewise, the contemporary workplace is saturated with technologies that alter the informational and material conditions of expert practice. Predictive algorithms, sensors, information systems, and cloud-based platforms are now standard features of knowledge-intensive work across nearly every industry. What is rarer, and what makes a context analytically generative for studying the intersection of expertise and technology, is a setting in which both conditions are present in particularly discrete and observable forms – where expertise has long been a defining feature of the work, where new technologies are visibly reshaping how that work is performed, and where the consequences of expert work are concrete enough to be seen.

Dairy farming offers precisely such a setting. The expertise of the average dairy farmer rests on knowledge that is largely tacit, intuitive, and embodied: the ability to read a cow's facial expressions, recognize early signs of mastitis by the look and feel of an udder, detect heat by the way a cow behaves, and integrate years of accumulated tacit knowledge into decisions about breeding, treatment, and culling. This is the kind of expertise that has

historically defined the practice of dairying and that experienced farmers describe in terms of having "a feel for cows." Over the past few decades, however, a generation of so-called informating and epistemic technologies has entered the dairy industry and begun to formalize, structure, quantify, and codify into data much of what was previously embodied in the farmer. Consequently, expert work is transforming in a multitude of ways, making dairy farming an unusually rich site for examining how technology participates in the (re)constitution of expertise.

The choice to focus on expertise emerged inductively from a larger comparative field study of dairy farms in California and New York that examined how emerging technologies and data were changing the practice of farming. My advisor invited me to work on this NSF-funded project and I was eager to put my experience working in large animal veterinary medicine to beneficial use. My work with ranchers, vaqueros, and pig farmers had instilled in me a sustained curiosity about the people doing the largely invisible and complex work of animal agriculture, and an appreciation for the expertise it takes to do that work effectively. Early farm visits surfaced a recurring set of concerns that informants returned to repeatedly: how technologies were changing what they needed to know, what they were able to do, and how their relationships to others shifted as new tools entered the farm. It was these initial observations that spawned this research focus.

A (Far Too) Brief Overview of Dairy Farming

Before explaining the research design and methodological approach of this dissertation, a brief overview of dairy farming is in order. The U.S. dairy industry occupies a substantial position in the agricultural economy. In 2024, U.S. farms produced approximately 226 billion pounds of milk from a national herd averaging 9.34 million cows, making milk one of the

country's largest agricultural commodities by volume (USDA National Agricultural Statistics Service, 2025). Beyond the farm gate, the broader dairy industry, including processing, manufacturing, and downstream sectors, generated an estimated $794 billion in total economic impact and supported roughly 3.2 million jobs as of 2023, according to industry-association analyses (IDFA, 2023). Milk itself is a federally regulated commodity. Most U.S. milk is sold under USDA Federal Milk Marketing Orders (FMMOs), which set minimum prices that processors must pay farmers based on formulas tied to USDA-surveyed wholesale prices of cheese, butter, nonfat dry milk, and dry whey (Greene, 2022; USDA AMS, 2023). Farm-level prices are further shaped by cooperatives, which market the large majority of U.S. milk on behalf of farmer-members and negotiate premiums above the regulated minimum (Whitt, 2024), and by global commodity dynamics that move the underlying wholesale prices on which FMMO formulas depend. Because milk is perishable, must be picked up daily, and is produced on farms that are small relative to the processors and retailers who buy from them, individual farmers function as price takers rather than price setters.

The fluid milk produced on farms moves through processors and cooperatives that turn it into the array of products consumers encounter on grocery store shelves (e.g., fluid milk, cheese, butter, yogurt, powdered milk, whey protein) with cheese and export-oriented manufactured products now accounting for a growing share of demand. Farmers operate within narrow and unpredictable margins determined by variable feed costs, labor availability, regulatory compliance (animal welfare, environmental, and food safety standards), and milk prices that fluctuate independently of their production decisions. The result is a perennially squeezed economic environment that has driven decades of structural change. Farms must now compete in volume. The number of U.S. farms selling milk fell 39% between the 2017 and

2022 Census of Agriculture, the largest decline between adjacent censuses since 1982, while total milk production rose 5%. By 2022, the 2,013 farms with 1,000 or more cows accounted for 66% of all U.S. milk sales. The industry, in short, is consolidating rapidly with fewer, larger, more capital-intensive operations producing more milk than ever before (Schmitz, 2025; Schnitkey, 2024).

At the industry level, dairy farming is best understood as a tightly networked ecosystem of specialized actors rather than a set of self-contained farms. A working dairy depends on regular inputs from veterinarians who oversee herd health and provide consultation on reproduction, husbandry, and disease prevention; nutritionists who formulate the feed rations that drive milk production; breeding specialists and genetics companies that supply semen, embryos, and artificial insemination services; hoof trimmers who manage lameness; equipment vendors and technicians who install and service milking systems, sensors, and software; feed and pharmaceutical suppliers; third-party labs that provide genomic and blood analysis; milk testing organizations that monitor quality and regulatory compliance; and milk haulers and cooperatives that move milk from farm to processor. Layered on top of this are agricultural lenders, insurance providers, regulators, and increasingly, technology vendors selling sensor-equipped devices and AI-enabled tools and data consultation services. These external experts visit farms on regular schedules and become deeply embedded in farm decision making, operational design, and physical and digital infrastructures, often working with the same farms for decades. The expertise required to run a modern dairy is therefore distributed across this network – farmers do not (and could not) do all of it themselves – and the work of running a farm involves continuously coordination between specialists, integrating their inputs, and translating their recommendations into operational decisions.

At the organizational level, dairy farms are typically family-owned businesses, often multigenerational, that have scaled into significant operations employing anywhere from a handful to several dozen workers. The labor force on most California dairy farms is heavily immigrant – primarily from Mexico and Latin America today, with earlier waves from Portugal and the Netherlands – and this demographic reality shapes neighboring dairy communities. Farms vary in their physical layout, herd size, and degree of technological integration, but they all share similar organizational structures and relational networks that are fine-tuned around milking and animal husbandry activities. In response to economic pressures, dairy farmers have been turning to technology to reduce costs while increasing milk production and cow welfare (Lage et al., 2024). Today, to manage and care for this extreme number of production animals, farms rely on a variety of tools to speed up, automate, support, eliminate, or improve work. This pressure to modernize and the rigid structures of farm work mean that technological and informational changes at the farm level (like the recent uptake of robotic milking systems) directly reshape the work of identifiable individuals rather than diffusing through layers of bureaucracy, which makes the dairy farm a particularly tractable site for studying how technology participates in expert work.

Activities Central to Dairy Operations

Although individual farms differ considerably in their technological configurations, scale, and management styles, the work performed on every dairy farm is organized around a common set of four key activities tied to the biological rhythms of the cow. Most visibly, milking structures the daily and weekly cadence of the farm. Cows are milked two or three times each day, every day, and the operation as a whole is calibrated to optimize both the volume and the quality of the milk that leaves the bulk tank. Surrounding milking is a

constellation of animal husbandry activities – detecting and treating mastitis, identifying and intervening in cases of illness, monitoring calf health for early signs of disease, ensuring cow comfort through proper housing, bedding, and hygiene, and managing feed delivery so that nutrition is matched to each animal's productive stage. A parallel reproduction program runs continuously alongside this husbandry work, encompassing heat detection, artificial insemination, pregnancy confirmation, and reconfirmation across the months of gestation, and the management of calving and post-partum care for both cow and calf. Finally, decisions about herd health and composition (e.g., which cows to cull, how to plan for replacement and population stability, how to maintain records, and how to coordinate with the veterinarian on herd-level health programs) provide the longer-horizon strategic frame within which the daily activities unfold.

These activities are deeply intertwined, and each one shapes and is shaped by the others. A cow's reproductive status determines her productive stage, which determines her nutritional needs, which influences her susceptibility to disease, which in turn affects her milk yield and quality, which feeds back into decisions about whether to breed her again or cull her from the herd. A mastitis case detected in the parlor triggers a treatment protocol that involves separating the cow from the herd, milk testing, interventions from herdsmen, the hospital pen manager, the veterinarian, and the records system. An undetected heat means a missed insemination, a longer interval to next calving, and a downstream loss of productive days that come with higher feed cost and lower value across the lifespan of the cow. Coordinating these interdependencies requires the continuous involvement of dozens of actors both on and off the farm – herdsmen, milkers, feeders, calf raisers, hospital workers, mid-level supervisors, functional managers, owners, and a rotating cast of veterinarians, nutritionists, breeders, hoof

trimmers, technologists, technicians, vendors, and consultants – all working in coordination toward the shared goal of supporting healthy, productive cows and farm profitability.

Expert work is therefore happening across the farm every day, organized through highly structured workflows and specialized roles whose outputs feed into one another. Decisions made by farmers and managers require deep and broad understanding of multiple intersecting domains simultaneously: the environment of weather, regulatory regimes, and political conditions; the industry dynamics of milk prices, labor markets, and disease outbreaks; the organization of farm operations and the coordination required to keep them running; the herd as an asset to be managed and planned for; and the individual cow with her particular disease processes, reproductive history, behavioral propensities, and nutritional requirements.

Expert Roles on the Farm

Roles on a dairy farm are organized around these activities and stratified by the breadth of decision-making authority each one grants. At the front line are task-based roles (e.g., milkers, bed rakers, feeders, pushers) whose work consists of executing a single critical task on a routine schedule with clearly defined goals, in close coordination with the other front-line roles whose tasks must align in time and sequence. Herdsmen and farmhands perform slightly more diversified sets of tasks, responsible for the physical management of cows and the execution of cow-related tasks that require more discretion than the single-task roles, such as identifying sick animals, administering treatments, moving cows between pens, performing breeding-related tasks, and assisting with calvings. Above the herdsmen sit mid-level supervisors who play a "jack of all trades" role, typically overseeing several herdsmen and assuming responsibility for one or two functional areas while supporting the herdsmen in performing more complex tasks. These roles may make perfunctory decisions about cow care

and herd management as needed. Functional managers are responsible for distinct sections of the herd or specific activity domains (e.g., heifer manager, milking parlor manager, robot manager, calf manager, feed manager, hospital manager) and operate as the operational specialists and decision makers for their respective domains. At the top, farm managers – typically one or two individuals, often a general manager and the farm owner sharing managerial responsibility – coordinate across functional domains and hold responsibility for the operation as a whole. Farm owners varied in their roles with some actively participating in day-to-day management of the farm while others solely focused on business and finances, depending on the operation. Smaller farms typically had flatter organizational hierarchies while maintaining the same decision-making accountability structure.

What holds this structure together is that although roles are highly stratified by decision-making authority, they are not siloed by function. Even at the most senior levels, farm owners may perform routine herd checks alongside their herdsmen or, in other operations, confine themselves entirely to business-level strategic decisions; functional managers typically know how to administer vaccines and identify mastitis even though those tasks fall to the workers they supervise; and milkers, whose formal role centers on the parlor, regularly detect mastitis cases, alert the hospital manager, and sort cows that need treatment. Each farm is structured as a kind of "all hands on deck" configuration in which everyone who works with cows is expected to know enough about cows to act on what they see. This cross-functional fluency reflects how careers in dairy work are typically built. Farmers usually begin in entry-level positions such as milking or feeding and develop their expertise through years of hands-on work with the animals, often asking to rotate into a herdsman role or being reassigned across farm functions as they accumulate experience. As they gain depth in successive areas (e.g.,

milking, hospital work, breeding, calf care) they may be promoted to lead positions, then to supervisors, then to functional manager, and eventually to general manager. The typical pattern of movement up the hierarchy and across functions means that expertise in this setting is prized and recognized precisely because it rests on this tacit and intuitive knowledge, demonstrable skills, and the established relational structures in which experts are embedded over years of "working their way up" in the farm. Strong expert identities are common in dairy work, and they are backed up by substantive skills and reinforced through recognition from peers. Expertise, therefore, is often observable, readable in the bodies of the animals, and durably remembered by a tight occupational community in which reputations, once established, tend to endure. Now that adequate context about how dairy work is structured has been laid, I now turn to the role technologies play in these core activities.

The Technological Evolution of the Dairy Industry

While farm activities have always been designed around cows, dairy technologies have been rapidly transforming the way these activities are performed. Over the past few decades, California dairy farms have steadily adopted interconnected suites of sensor-based and AI-enabled systems, including automated robotic milkers, in-line milk quality analyzers, activity tracking collars, RFID identification tags, genomic testing services, predictive algorithms, and integrated herd management software platforms. Together, these technologies have produced vast quantities of previously unavailable data about individual cows and generate algorithmic predictions that increasingly participate in decisions farmers historically made through embodied judgment. Detecting mastitis, identifying cows in heat, predicting illness, and ranking animals for culling and breeding – all activities that traditionally relied on specialized

knowledge and hands-on examination – are now mediated by devices and software that produce continuous, granular, individualized data and offer recommendations.

These technologies have been gradually and incrementally improving over years as farms continuously upgrade their infrastructure. However, informants in this study noted four major inflection points over the past three decades that reshaped how they experienced dairy work. The earliest of these was the RFID tag revolution in the early 2000s (see **Figure 2.1** for example), which replaced paper-based record keeping with electronic identification ear tags integrated into digital databases. Where farmers and herd managers had previously recorded individual cow data by hand on physical cow cards or entered it manually into databases at the end of the day, RFID-enabled handheld devices and platform-based systems allowed real-time data capture at the point of interaction with each animal. Informants reported that workflows became dramatically faster and less error-prone, and the digitization of individual cow records laid the infrastructural foundation on which subsequent technologies would build.



Figure 2.1. Image of an RFID ear tag (grey button tag on the left ear) on a Holstien calf. Tags are applied at birth and remain on the cow throughout her entire life, providing a digital form of ID that connects to a digitized record.

A parallel transformation in the mid-2000s occurred in the shift from conventional milking parlors which captured milk volume in bulk, to rotary milking parlors equipped with per-stall sensors which enabled milk production to be tracked at the level of the individual cow. Rotaries (see **Figure 2.2** for example) also accelerated throughput by moving cows around a carousel while workers performed multiple tasks concurrently, enabling farms to milk substantially more cows per hour while generating finer-grained data on individual cow health and productivity.



Figure 2.2. An image of a rotary milking parlor taken from the DeLaval website.

Two further technological transformations followed, each extending the reach of granular, individualized data into new domains of expert judgment. The current generation of activity tracking collars (see **Figure 2.3** for example) that equipped cows with accelerometers that captured movement, feeding, and rumination patterns around the clock began gaining traction in the mid-2010s, providing data streams that could identify estrus behavior earlier and more accurately than visual heat detection alone. This increased the precision of reproduction programs that had previously depended solely on the trained eye of the breeder.



Figure 2.3. An image of a Holstein cow wearing an activity tracking collar taken from the AllFlex website.

Most recently in the late 2010s, robotic milking systems automated the milking role itself while generating real-time, per-cow, per-udder-quarter data on yield and milk quality (see **Figure 2.4** for example). The most recent generation of robotic milking systems incorporates in – line milk chemistry analyzers capable of detecting subclinical mastitis and automatically diverting contaminated milk before it enters the bulk tank, shifting quality control from retrospective testing to continuous monitoring.



Figure 2.4. An image of a cow being milked in a Lely A5 Astronaut robotic milker taken from the Lely website.

Running alongside these on-farm transformations, genomic testing became widespread in the late 2000s and has since transformed breeding decisions by enabling prospective predictions about an animal's health, conformation, and productivity at birth or even as an embryo (see **Figure 2.5** for an image). Where generational improvement in dairy genetics had once depended on observing an animal's lifetime performance and that of its offspring, genomic models built on decades of accumulated genetic and phenotypic data now allow farmers to iterate on artificial selection across generations at unprecedented speed, exercising new forms of control over hereditary traits that previous generations of farmers could approach only through observation and iteration.

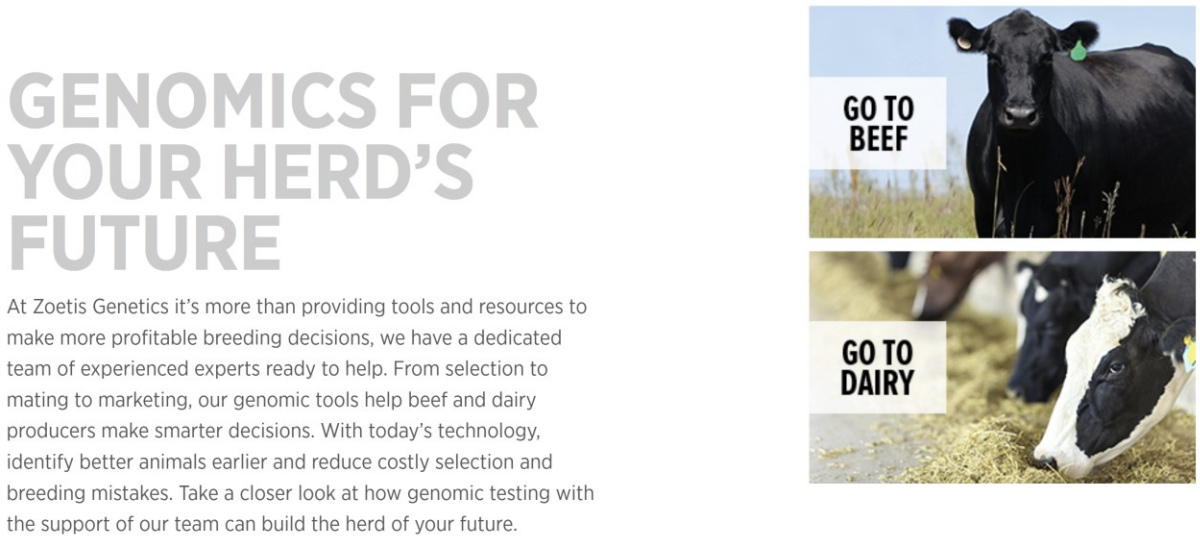


Figure 2.5. A screenshot taken from the Zoetis Genomics website showing the marketed use case for genomics testing for breeding decision support.

The Modern Dairy Farm

Modern dairy farms operate through a web of interconnected devices, platforms, and systems that together quantify, store, analyze, and act upon information about every cow in the herd. Although the specific configurations vary across farms – each operation assembles its own combination of platforms, vendors, and hardware – these systems can be grouped into four functional categories that interact continuously throughout the day. Sensing systems, including milking robots, activity tracking collars, in-line milk meters, and RFID scanners, do the foundational work of translating physical phenomena into digital information. They detect, informate, and digitize the embodied behaviors and biological signals of individual cows, producing the digital representations on which all subsequent work depends. The data these systems generate flow into data reporting systems – herd management software platforms such as DairyComp and DHI-Plus that are accessible through computers and phone apps, and robot-specific platforms such as Lely's Horizon and DeLaval's DelPro – which function as the

primary repositories where individual cow records accumulate, are accessed, and are reviewed by farmers, veterinarians, and other contractors. These platforms are the day-to-day reference points around which farm decision making is organized – the digital cow cards of the contemporary dairy.

Atop this data infrastructure sit two further classes of technology that turn stored information into actionable metrics. Analytic systems, such as AllFlex's heat detection module and the predictive algorithms embedded in Horizon and similar platforms, process raw sensor data into interpretive outputs that farmers can use to make decisions (e.g., predictive scores for mastitis risk, heat probability, and illness; binary condition flags like "due for vaccination"; and categorical rankings that sort cows into groups based on dozens of differentiating factors). These analytic outputs configure attention, telling farmers and herdsmen which cows to look at, in what order, and for what purpose. Finally, automation systems perform physical work directly. Robotic milkers attach to cows and extract milk without human intervention, automated sort gates separate cows into hospital, breeding, or fresh pens based on algorithmic triggers, and in-line milk chemistry analyzers detect contamination and automatically divert affected milk from the bulk tank. As the data map in **Figure 2.6** illustrates, the data generated on a modern farm circulate across all four of these system categories and across the human and organizational actors who depend on them (e.g., farmers, herdsmen, veterinarians, breeders, nutritionists, vendors, and the technology companies whose platforms host the data) such that information captured at the moment a single cow enters a milking stall may, within minutes, inform a veterinarian's treatment recommendation, a breeder's insemination decision, and a herdsman's morning hospitalization list. The result is a tightly entangled technological web in

which expert work is performed not in isolation but in continuous dialogue with the digital traces of the herd.

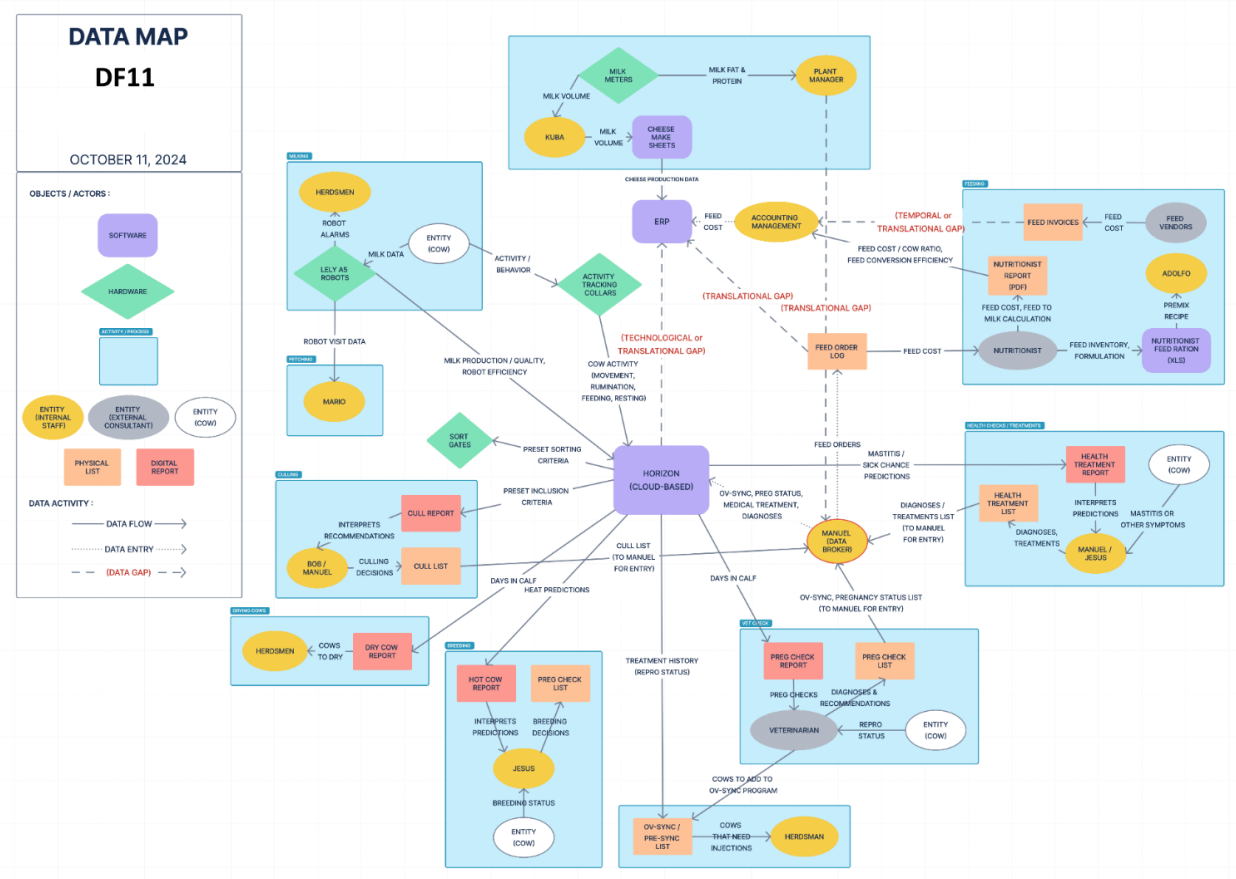


Figure 2.6. An example of a data map from DF11 depicting all the data-generating devices and the linkages between human, technological, and physical entities through which data flows.

By contrast, more conventional dairy farms (what informants described as “low-tech” farms) operate with a far thinner technological substrate, one in which much of the information that flows automatically through sensors and platforms on a modern farm must instead be generated, retained, and transmitted through the bodies and memories of the people who work there. On a conventional farm, milking occurs in a parallel parlor where milk is measured in bulk rather than per cow, foreclosing the individual-level production data that drives so much decision making elsewhere. Cows are identified by visual numeric ear tags rather than RFID,

which means a herdsman who needs to act on a specific cow must first locate her by eye and read her number. Detecting issues like pneumonia depends on familiarity with the herd, attentiveness to subtle sensory cues, and the accumulated knowing that experienced dairymen describe as "having an eye for cows." Records are kept on paper cow cards or entered by hand into databases at the end of a shift, introducing both delay and the possibility of data entry error. Heat detection is performed visually, often at dawn when behavioral signs are most apparent, and depends on a breeder's vigilance when walking pens to visually scan for physical proxies such as evidence of mounting indicated by rubbed off chalk marks that he must reapply daily. Illness is identified through the embodied diagnostic practices that have long defined dairy expertise (e.g., watching the cows, noting the appearance of manure, palpating udders, looking at the color and texture of milk, listening for coughing). External experts such as veterinarians visiting these farms work without digital insights, relying instead on the farmer's verbal account, their own observations, and the relational history they have built with farms.

The contrast I draw here is deliberately schematic. It aggregates descriptions of "low-tech farms" from farmers and veterinarians of the data-thin end of the spectrum opposite the data-intensive extreme. In practice, nearly every contemporary dairy farm has adopted some form of digital infrastructure, but the specific tools, the depth of their integration, and the combinations in which they are deployed vary widely from farm to farm – no two operations assemble this panoply of technologies in quite the same way. The conventional farm I describe here is therefore best understood as an analytic ideal type, the opposite pole against which the data-integrated farm can be contrasted, rather than as a portrait of any one site in my sample. The table in **Appendix A** summarizes the key technologies, techniques, and processes that

differ between conventional and data-intensive dairy farm operations across the three studies in this dissertation.

Dairy Farming as an Ideal Research Context

Dairy farming is a sophisticated form of expert work that makes it unusually generative for studying technology's participation in the constitution of expertise. Successful farmers must continuously monitor herds of hundreds to thousands of animals while making consequential decisions about health interventions, breeding, feeding, and herd management that directly affect both animal welfare and the financial viability of the operation. This work draws on a body of knowledge spanning cow biology, behavior, psychology, nutrition, genetics, reproduction, disease management, and business operations – expertise that has traditionally been developed through years of hands-on experience and embodied in sensory skills. Experienced farmers can typically detect subtle changes in cow behavior, identify signs of illness through visual and tactile examination, and make intuitive judgments about animal health and productivity based on accumulated pattern recognition through what veterinarians and farmers describe as "cow sense." This expertise is at once technical and relational: farmers construct their identities as experts through ongoing interactions with animals, with their own employees, and with a dense ecology of external specialists – veterinarians, nutritionists, breeders, hoof trimmers, technicians, and consultants – who participate continuously in the work of running a farm.

What made dairy farming particularly generative for this dissertation was that it rendered all three constitutive processes of expertise (knowing, performing, and recognizing) simultaneously visible within a single empirical setting. Most studies of technology and expert work focus on one of these processes at a time, in part because most settings make only one

accessible. In a dairy context, all three were observable at close range within bounded organizational and relational structures that allowed the cascades among them to be traced. The knowing dimension was observable because farmers' epistemic relationship with cows were both deep and visible. Cows were the central object of farmers' expertise, encountered daily through embodied means (e.g., sight, smell, sound, touch) and became increasingly mediated through data and digital representations. The transition from sensory to data-based ways of knowing cows was an ongoing transformation that could be observed in real time as farmers moved between the barn and the screen. The performing dimension was observable because farmers had to continuously demonstrate their expertise to a dense ecology of audiences embedded in their daily work. Coworkers, vendors, veterinarians, breeders, nutritionists, and consultants were present on the farm routinely, and farmers had to adjust how they presented their expertise when new technologies disrupted established performance. The recognizing dimension was observable because farms operated within tightly networked, long-duration relationships with external experts who repeatedly evaluated farmer competence over time. A veterinarian may have served the same farm for twenty or thirty years. These sustained collaborations created stable, observable evaluation processes that carried real consequences for how farmers were positioned within their expert work networks. Because all three dimensions were simultaneously visible within the same setting, dairy farming allowed me to trace not only how technology participated in each constitutive process but also how changes in one spilled over into changes in the others.

Beyond this theoretical fit, dairy farming offered several practical methodological affordances. First, the full life cycle of representational work was encompassed within a single organization. The physical referents (e.g., cows, milk), the technologies that converted them

into data (e.g., robotic milkers, activity collars, in-line analyzers, lab testing), the computational systems that processed data into representations (e.g., herd management software, predictive algorithms, composite scoring models), and the users who interpreted and acted upon those representations (e.g., farmers and their external specialists) were all co-located on the farm. A researcher could follow the entire process – from observing how data was extracted from a cow, to how farmers engaged with different snapshots of representational data, to how they acted upon those representations – within a single day. Second, cows themselves were productive empirical referents because they occupied an intermediate ontological status. They could be bred for specific traits, treated for illness, trained to conform to behavioral expectations, and culled from the herd in ways that cumulatively reshaped bodies, productive capacities, and herd composition; yet they also resisted change, developing according to biological rhythms and dispositions that were not fully programmable. This intermediate status made the material consequences of representational evaluation concrete and observable while ensuring that the work required to align representations with physical realities remained visible. Third, the organizational structure of dairy farms supported the kind of analysis this dissertation required. As described in the previous section, farms typically operate with relatively flat structures and concentrated, operationally proximate decision-making authority, which meant technological and informational changes directly reshaped the work of identifiable individuals rather than diffusing through bureaucratic layers. Tracing how technology participated in expert work was therefore tractable in a way it would not have been in a more bureaucratically complex organization. Finally, the tight ecology of on-farm and external specialists who visited farms on regular weekly or biweekly schedules meant that

audience-expert dynamics unfold at close range and on predictable cycles, making them observable through routine fieldwork rather than through one-off encounters.

Research Design and Epistemological Approach

This dissertation adopts a relational ethnographic approach grounded in inductive, interpretive epistemology. Relational ethnography, as articulated by Bailey, Faraj, Hinds, Leonardi, and von Krogh (2022), focuses analytic attention on the relations between entities (e.g., humans, technologies, animals, documents, organizations) rather than on the properties of isolated actors. To capture co-constitution, I needed a methodological approach that does not treat technology as an external variable acting upon expert work, but rather as an entity whose participation in relational dynamics co-constitutes it. Relational ethnography also enables the analysis to move across the three constitutive processes, tracing how expert-object relations (knowing) cascade into expert-audience interactions (performing) and into audience-expert evaluations (recognizing). This approach also draws on the broader tradition of workplace ethnography that takes seriously the notion of "bringing work back in" to understand how work is accomplished through embodied engagement with materials, tools, and other people (Barley & Kunda, 2001).

The research questions animating this dissertation concern processes – how experts come to know, how they perform, and how they are recognized – that unfold *in situ*, across time, and through interaction. They require methods capable of tracing how practices develop, how interpretations shift, and how relationships among humans, technologies, and cows evolve. Expertise itself involves tacit, embodied, and relational dimensions that informants cannot fully articulate. Farmers know things about their cows that they cannot easily put into words; veterinarians develop inferences about farmers that they would not necessarily

volunteer; and the way a farmer toggles between a screen and a cow during a routine task reveals interpretive work that can only be characterized through observation. Further, interactions with technology can be notoriously difficult to study, even qualitatively (Bailey et al., 2021). Ethnographic methods are therefore most suitable for the emergent nature of the phenomenon under study (Edmondson & Mcmanus, 2007).

Site Selection

I sought farms at every stage of technological integration with varying combinations of the four categories of informating and epistemic technologies described above to provide sufficient operational and epistemic variety to surface knowing, performing, and recognizing dynamics. I held constant the institutional context by drawing all sites from California, where market, regulatory, economic, and geographic conditions are comparable across operations. A summary of field sites appears in **Appendix B**.

Each of the three studies in this dissertation draws a purposive sub-sample from the dataset, selected to maximize analytic leverage on its specific research question. Study 1 draws on six farms employing the full data-intensive technology suite (genomic testing, activity tracking collars, and automated robotic milkers paired with predictive algorithms and herd management software). Study 2 draws on four California farms that implemented virtually identical robotic milking systems within similar time periods. Study 3 draws on twelve California dairy veterinarians and the seventeen dairy organizations they served. The detailed selection rationale and inclusion criteria for each sub-sample appear in the relevant individual study chapters.

Data Collection

Participant observation enabled me to immerse myself in the work practices and interactions structured around farm work. Accordingly, I designed this study to prioritize depth, breadth of variation, and access over total observation hours. Rather than conducting long-duration immersion at one or two sites, I structured fieldwork around the natural cyclicality of dairy operations. Because routine farm activities are organized around week-long cycles encompassing daily milking, breeding, herd checks, vet visits, feeding, calf management, and machine maintenance, week-long intensive observation at each focal site captured the full range of routine operations.

Data was collected over three years, between August 2022 and November 2025, and unfolded in three phases. Phase one involved visits to robotic farms to learn about the activities and actors involved in milk production, conducted between April and July of 2023 at DF6 and DF9 (with earlier data collected from DF2 spawning this dissertation research). Phase two involved more targeted data collection examining automated milking systems and data-intensive technology use through participant observation at DF10, DF11, and DF12 between May and August 2024, along with follow-up interviews with members of DF1, DF6, and DF9. Phase three involved shadowing large animal veterinarians in November 2025, with observations and interviews of client farmers occurring in the process of participant observation with veterinarians. In total, I spent seven weeks conducting participant observation across five focal dairy farms and two weeks shadowing veterinarians. In total, I visited, interviewed, and observed activities at sixteen dairy farms and two calf ranches.

Field observations were conducted across all major farm activities, including milking, feeding, breeding, calving, hospital and medical management work, herd checks, vet checks,

routine treatments, and office and decision-making work. Observation sessions typically lasted two to eight hours, capturing not only formal work but also unstructured moments (e.g., car rides between farms, lunch breaks, debrief conversations, informal exchanges) in which informants candidly reflected on their work and their counterparts. I kept shorthand notes in a notepad during observations unless rain made it impractical. Informal interviews were conducted *in situ* while shadowing informants, allowing reflection to remain anchored to specific cases and decisions (Van Maanen, 2011). Semi-structured interviews lasted between thirty minutes and two hours and were conducted where deeper accounts were needed (Spradley, 1979), particularly to establish pre-implementation baselines in Study 2. All interviews and observations were audio recorded and transcribed with the help of Otter.ai and/or by hand. Photographs and videos documented software interfaces, dashboards, paper records, physical work settings, and complex work processes. Artifacts collected included paper reports, protocols, standard operating procedures, screenshots of algorithmic outputs, treatment lists, handwritten notes, and benchmarking spreadsheets shared between veterinarians and farmers. Between site visits, I maintained infrequent contact with farmers through text messages and email, allowing me to track developments while unable to be in the field (e.g., during the 2024 H5N1 bird flu outbreak when farms were in lockdown due to infections spreading across CA dairy herds), ask follow-up questions, and observe how interpretations of events evolved over time. Supplementary data included publicly available media (e.g., trade publications, web articles, YouTube interviews, blog posts, and news coverage of farms and the technologies they employ) which contextualized observed practices within broader industry discourse.

In total, I observed, interviewed, and interacted with seventy-four farm workers across sixteen dairy farms and two calf ranches, as well as twelve veterinarians across five large animal veterinary practices, and conducted sixty-six semi-structured and informal interviews. See **Appendix C** for a full list of informants. I also interacted with dozens of external vendors, contractors, technology representatives, breeders, nutritionists, academics, and service providers during the project whose individual identification is not tabulated here. A summary of the data corpus appears in **Table 2.1**.

Table 2.1. Summary of collected data.

Organization	Time spent in the field	Number of Informants	Number of interviews	Photos, videos, docs
DF1	–	1	2	–
DF2*	–	–	–	–
DF6	1 week	20	10	372
DF9	1 week	12	5	152
DF10	1 week	6	2	79
DF11	2 weeks	10	12	121
DF12	1 week	5	4	61
Other farms (11)^	–	20	18	343
Veterinarians	2 weeks	12	13	75
Total	**9 weeks**	**86**	**66**	**1203**

* Data collected from DF2 were drawn from the dataset for a broader pilot study that spawned this dissertation research. Activities represented in this table only include data personally collected by the author.
^ "Other farms" denote aggregate metrics for farms visited and observed for one day, including interviews with farm members as part of veterinary herd checks, not farms included as field sites.

Concretely, participant observation allowed me to appreciate the physicality of dairy work in ways that shaped both the data I collected and how I came to interpret it. Each day, I woke up around 3am to arrive at the farm for the morning shift. Milking happened around the clock, so farms were always bustling with activity no matter the time of day or night. Farmers worked during freezing temperatures, 110°F heat, scorching sun, pouring rain, pandemics,

swarms of biting flies, ankle deep in excrement for hours each day, surrounded by hundreds of half-ton reactionary prey animals. I was lucky enough to get a front-row (and backstage) view of the work, which was often physically demanding, dirty, frustrating, and unpredictable (see **Appendix D** for a selection of images from field work). Workers often had to corral these gentle giants into narrow stanchions so they could be restrained, examined, palpated, treated, or inseminated. Domesticated cows are not aggressive creatures by nature; however, they are massive and surprisingly agile so getting one to move requires the right combination of strength, finesse, instinct, and nerves. Standing in the direct path of a running cow to stop her in her tracks is not an action typically written about in ethnographic field notes, nor is the act of shoving one's arm up a cow's rectum to scan for fetal heartbeats with a portable ultrasound. And yet my notes are rife with descriptions of these work practices that take place on the farms every day which I tried to record faithfully to capture as much of an emic perspective on the realities of farm work as possible (Geertz, 1973).

Data Analysis

Analysis across the three studies followed an inductive, iterative approach consistent with grounded theorizing (Corbin & Strauss, 2015; Glaser & Strauss, 1967). Across each study, I moved back and forth between the data and relevant theoretical literatures, allowing patterns to emerge from the data while remaining attentive to how those patterns engaged existing scholarly conversations. The data corpora and analytical procedures for each study were tailored to its specific research question while sharing the dissertation's common inductive, grounded orientation. The detailed accounts of each study's methods appear in the relevant chapters. Here, I summarize the data each study draws on and the analytical move that defined it.

Study 1: Knowing as a Recursive Process of Representing – How Data-Intensive Technologies Reconstitute the "Good Cow"

Study 1 examines how data and predictive technologies affect what experts know about their objects of expertise and how they come to know them. Answering this question required data capable of capturing the moment-to-moment movement between physical cows and digital representations. Observational data captured this movement directly as researchers watched farmers shift attention between the screen and the animal, weigh sensor outputs against embodied observation, and translate representational judgments into physical interventions. Informal interviews conducted while farmers engaged with the technology produced grounded, concrete articulations of how data was interpreted, how it was weighed against what farmers could see and feel, and how it was translated into action. Artifact collection (e.g., screenshots, reports, algorithmic outputs) preserved the digital representations themselves, allowing analytic reconstruction of what farmers were knowing at each moment of decision. Week-long visits across the full operational cycle ensured that the full range of representational moments was captured: the inscription points where cows became data, the biographical records that accumulated over time, the algorithmic predictions that converted data into projected futures, and the comparative rankings that situated cows within the herd. The data corpus assembled for Study 1 is appropriate to its research question because it captures the entire life cycle of representational work as it unfolds in farmers' epistemic relationship with cows.

Analysis began with open coding of activities, interactions, and practices across the field sites and then moved toward coding for farmers' evaluative judgments about cows after I identified a pervasive linguistic pattern in which farmers used "good cow" and "bad cow" to mean different things in different representational contexts. Successive rounds of coding

deconstructed the representing process into four renderings – inscription, biographical, algorithmic, and comparative – each restructuring how cows were evaluated and prompting different forms of congruence work between digital representations and the physical animals to which they referred.

Study 2: Performing as a Recursive Process of Orienting – How Pre-Existing Performance Threats Shape Technology Use in Practice

Study 2 examines how farmers adapt their performance of expertise in response to the threat that AI-enabled robotic milking systems pose to their established expert standing. Observational data of farmers and their interactions with others captured these performances directly. Comparison across four nearly identical farms allowed observed variation in performative response to be attributed to farmer-level adaptation rather than to differences in technology, market, regulatory, or geographic context. Retrospective accounts gathered through semi-structured interviews provided a comparative baseline against which post-implementation performances could be assessed – what farmers had done before the robotic system arrived, how they had presented their expertise to the same audiences, and how this had changed. A single observation could capture an instance of a farmer scrutinizing an algorithmic output; sustained observation captured the consistency of orientation across days, audiences, and decision contexts that allowed me to identify stable adaptive strategies rather than isolated reactions. The data corpus assembled for Study 2 is appropriate to its research question because it captures performances of expertise as recurring, audience-directed activity.

Analysis moved from open coding to selective coding for farmers' practices in the presence of robotic milking systems and developed a typology of orientations toward the technology, with cross-farm comparison revealing how different orientations led to different

performance strategies through which farmers managed the threat the technology posed to their expert standing.

Study 3: Recognizing as a Recursive Process of Inferring – How Audiences Construct Expertise Inferences Under Digital Visibility

Study 3 examines how technologies shape how veterinarians, as expert audiences, evaluate and recognize farmer expertise. Answering this question required data from actors with comparative perspective on farmer competence across many farms and across extended time horizons. Veterinarians occupy this structural position uniquely. They serve dozens of farms simultaneously, often over decades, and they encounter both data-intensive and conventional clients within their weekly rounds. Observation during vet visits captured both the formal evaluative encounter (e.g., herd checks, pregnancy confirmations, treatment decisions) and the unstructured moments where candid evaluations of farmers emerged: car rides between farms, lunch breaks, debrief conversations among partners. These unstructured moments were where vets compared farmers explicitly, narrated the reasoning behind their judgments, and articulated the inferences they drew from the digital traces of farmers' backstage work. Semi-structured interviews allowed vets to articulate their inferences about farmer expertise more reflectively and to discuss how digital tools had altered those inferences over time. My prior background in veterinary medicine facilitated trust and access to candid accounts that would have been buffered with an outsider. The sample of twelve vets, representing approximately six to twelve percent of California's active dairy practitioner population, provides meaningful coverage of a small and specialized professional community. The data corpus assembled for Study 3 is appropriate to its research question because it captures recognition as a sustained, comparative, and increasingly digitally mediated process.

Analysis focused on selective coding across veterinarian interviews and field notes and revealed a consistent pattern in how vets explicitly evaluated farmers ("good cow guy," "rock star," "no cow sense," "wild card"). This pattern prompted theorizing about how veterinarians constructed inferences about farmer expertise from the digital traces of backstage work.

CHAPTER 3: Knowing as a Recursive Process of Representing – How Data-Intensive Technologies Reconstitute the "Good Cow"

Digital representations – patterned sets of digitized symbols that stand in for physical phenomena – have become pervasive infrastructures through which organizations know, evaluate, and act upon the world (Bailey et al., 2012; Kallinikos, 1995; Kallinikos et al., 2013; Zuboff, 1988). From automotive crash simulations that substitute for physical prototypes (Leonardi, 2012) to algorithmic systems that classify and rank individuals at a distance (Bowker & Star, 1999; Brayne, 2017), digital representations increasingly mediate the relationship between organizational actors and the physical realities they seek to understand and manage. This mediating role has prompted sustained scholarly attention to the consequences that digital representation have for reshaping work and expertise (Barley, 1996), enabling new forms of comparison and evaluation (Espeland & Stevens, 1998, 2008; Muniesa, 2011) and decoupling organizational knowledge from the physical phenomena it purports to describe (Borgmann, 1999; Burton-Jones & Grange, 2013). These studies have established that digital representations are consequential for how organizations know and evaluate the world, demonstrating that the transition from direct observation to digitally mediated knowing transforms not only what is visible; but also what is valued, what is acted upon, and how work itself is organized around representations (Baskerville et al., 2020; Østerlie & Monteiro, 2020). Yet as representational technologies grow more continuous, granular, predictive, and autonomous an important question demands renewed attention: How, precisely, do digital representations come to reshape not only what we know about the physical world, but also the physical world itself?

Literature in organizations and information systems has examined what happens when organizational actors work *with* representations after they have been created (Alaimo &

Kallinikos, 2022; Recker et al., 2021). This literature has produced important insights about the consequences of representational use, from the simulations that "lure" people into thinking they no longer need to work with physical referents (Bailey et al., 2012) to the evaluative apparatuses that determine what counts as visible and valuable (Espeland & Stevens, 1998; Orlikowski & Scott, 2014). But these studies treat representations as 'things' that individuals work *with* as standalone entities. Consequently, far less attention has been devoted to sustained, multi – layered work through which physical phenomena are inscribed, encoded, rendered, and continuously reconciled with the digital forms that stand in for them. This organizational process of *representing* itself shapes organizational realities. While the science and technology studies (STS) tradition has long recognized that representing is constitutive rather than merely descriptive (Goodwin, 1995; Latour & Woolgar, 1986; Mol, 2000), organization and information systems (IS) scholarship has largely treated representations as artifacts, attending to what they *do* once they exist rather than how they come into being and what that process does to the objects being represented.

The distinction between working *with* representations and the work *of* representing is consequential because much of the organizationally meaningful action – the establishment of evaluative criteria, the framing of objects, and the interventions these provoke on physical referents – occurs *during* the work of representing rather than after representations have stabilized into artifacts. This paper explores "representing" as an organizational process in its own right; one that involves continuous work to achieve and sustain faithful correspondence between digital representations and their ever-changing physical referents. Drawing on an ethnographic study of data-intensive technologies and complex representational systems in dairy farming, I examine how the organizational process of representing exerts distinct yet

cumulative selection pressures on physical referents (i.e., cows). Crucially, this process does not produce a single, monolithic depiction of the represented entity. Rather, different phases and practices of representing generate distinct *renderings* – partial, distinctly framed versions of the same digitalized object that each make different aspects visible, apply different evaluative criteria, and present the object as a different kind of thing. The same physical entity might appear simultaneously as a body to be transformed into data, a biographical record of past events, an algorithmically projected future, or a position within a ranked distribution. Each rendering offers a snapshot or glimpse into a complex representation, but each framing what the object is in ways that invoke different questions about its worth. I demonstrate how attempts to faithfully capture physical reality ended up (directly and indirectly) reshaping the physical referent through a highly entangled process of redefining the "goodness" of cows and driving farmer efforts that physically transformed them over time.

As organizations pursue ever more ambitious digitalization efforts – digital twins, remote sensing, predictive health monitoring, real-time performance tracking – these practices of representing operate continuously, at scale, and with growing autonomy. Under these conditions, the most consequential organizational dynamics may reside in the process of representation itself. What matters is not simply what representations show actors after the fact, but how the ongoing work of rendering physical phenomena into digital forms restructures what those phenomena are, what they are worth, and what they become.

Digital Representations and the Physical World

Composed of data objects such as datasets, images, metrics, and model outputs, digital representations extend a long semiotic tradition in which signs stand in for objects through referential symbolism (Burton-Jones & Grange, 2013; Kallinikos et al., 2013). What

distinguishes them from their analog predecessors is a set of properties intrinsic to computational encoding – editability, interactivity, reprogrammability, openness, and distributability – that allow them to circulate across contexts, be endlessly recombined, and serve as the underlying digital matter that shapes organizational reality (Alaimo & Kallinikos, 2022; Kallinikos et al., 2010; Østerlie & Monteiro, 2020). As these representations have proliferated, they have come to mediate physical experimentation (Bailey et al., 2012; Leonardi, 2012), shape categories and interactions (Alaimo & Kallinikos, 2021; Bowker & Star, 1999), enable comparison across qualitatively disparate domains (Espeland & Stevens, 1998; Orlikowski & Scott, 2014), and increasingly substitute for direct observation of the phenomena they encode (Brayne, 2017; Zuboff, 1988).

Working *With* Representations

The prevailing approach to digital representations in information systems and organization studies treats them as artifacts that have already taken shape by the time they reach analytical or practical utility. Whether the artifact in question is a crash simulation that engineers consult in place of physical prototypes (Bailey et al., 2012), an algorithmic ranking that reorders competition among hotels (Orlikowski & Scott, 2014), or a screen-based display that mediates operators' contact with a pulp mill (Zuboff, 1988), the analytical starting point is a stabilized representation whose effects on organizing can then be traced. Thus, the central research question in this tradition concerns how organizational actors work *with* representations and the consequences that follow. Under this framing, scholars ask what representations do once they exist: how they structure cognition, coordinate action, redistribute authority, and risk supplanting the physical realities they depict (Alaimo & Kallinikos, 2022;

Bailey et al., 2012; Baskerville et al., 2020; Burton-Jones & Grange, 2013; Recker et al., 2021; Weber, 2003).

This body of work has produced important insights. Working with digital representations can alter how people understand the phenomena being represented, producing a "lure of the virtual" (Bailey et al., 2012) or a "dangerous illusion of the connection between data and reality" (Boyd, 2021, p. 549) in which representations are mistaken for faithful substitutions of physical referents. It can profoundly alter the experience of work, as Zuboff's (1988) seminal study of paper mill operators demonstrated when digital mediations of chemical and mechanical processes – previously monitored by sight, smell, and feel – required workers to develop new intellective skills as representations began to "replace concrete reality" (p.63) and the "sense of the hands-on" (p.65). And it can reshape evaluation by rendering unstructured phenomena categorizable, comparable, rankable, and evaluable through mechanisms such as selective objectification (Kallinikos, 1995), commensuration (Espeland & Stevens, 1998), and the evaluative apparatuses that "define and include some things in certain ways, precluding and excluding others" (Orlikowski & Scott, 2014). Even Porter (1995, p. 11) reminds us that "the credibility of numbers…is a social and moral problem," revealing the social underpinnings of what comes to count as accurate or authoritative within representations.

Across these studies, the dominant concern is representational decoupling: the risk that digital representations will supplant or 'break free' from their physical referents as the primary object with which we work. Burton-Jones and Grange (2013), quoting Weber (2003), articulate the foundational premise: "The raison d'etre for information systems [is] that they track states of and state changes in other systems. By observing the behavior of an information system, we

obviate the need to observe the behavior of the system it represents" (p.636). Others warn that digital representations have "the greatest potential to change work's historically tight coupling to the physical and, with it, the work relations of people to objects and each other" (Bailey et al., 2012, p. 1486). Taken to the extreme, digital representations can lead to "liquefaction" (Monteiro & Parmiggiani, 2019, p. 167) or an "ontological reversal" (Baskerville et al., 2020, p. 509) between the digital and physical, a prospect Baudrillard (Baudrillard, 1994) captured in his observation that "the territory no longer precedes the map, nor does it survive it" (p.1-2).

Yet the evidence suggests that this decoupling rarely occurs. Empirical studies consistently show that when confronted with incongruence, individuals respond by shoring up the linkage between digital and physical rather than abandoning the physical altogether. Zuboff's (1988) mill operators continued to monitor pulp status through sensory means to "conquer the felt distance of the referential function" (p.65). Bailey and colleagues' (2012) crash test engineers returned to physical cars to verify what they observed digitally. Østerlie and Monteiro (2020) demonstrated how laboratory assistants went to considerable lengths to incorporate representations into everyday practices that made them "organizationally real" (p.1). In each case, actors worked toward representational congruence – a form of continuous stewardship that Barley (1996, p. 419) refers to as "caretaking" to ensure representational and physical systems remained intact and in "good working order" to sustain faithful correlation to the physical referent such that working within the representation remains useful (Weber, 2003). A "good" representation, as Burton-Jones and Grange (2013, p.636) argue, must "faithfully represent some domain, because [faithful representations] provide a more informed basis for action than unfaithful representations do."

This persistent pursuit of congruence between representations and their physical referents reveals an important boundary of the working-with-representations literature. If actors are continuously working to maintain the tight coupling between digital and physical that existing literature cautions might be severed by placing too much faith in digital representations, then the interesting question is not only what representations do to organizing after they are formed, but what happens during the organizational process through which that congruence is pursued and maintained. This points toward a fundamentally different set of questions about what effects representing has on organizing that prior studies allude to but do not fully engage.

Working *Towards* Representation

Acknowledging that representations require continuous congruence work raises a prior question: how do physical phenomena become digitally represented in the first place? While the organization and IS literature has focused predominantly on what happens after representations take shape, the science and technology studies (STS) tradition has long attended to the constitutive work involved in constructing representations. Across this tradition, three insights stand out that together reframe representing as a multi-layered process with significant evaluative consequences.

The first insight is that representing is not a neutral act of encoding pre-existing reality. The instruments, conventions, infrastructure, and practices through which physical phenomena are rendered into signs, symbols, and data actively constitute the objects they purport to merely describe. Latour and Woolgar (1986) demonstrated this through their concept of inscription devices, showing that laboratory instruments do not simply record but actively produce the phenomena they measure. Goodwin's (1995) study of scientists mapping the Amazon River

shows how the tools available for quantifying depth, current, and sea-floor features determined what could count as the river as an object of inquiry. Pinch (1996) extends this argument to particle physics, where what counts as a "particle" is itself a product of social and technological conventions embedded in instruments. Lynch (1991) illustrates how scientists transform invisible phenomena and noisy data into diagrams, graphs, and visual representations through complex instrumental mediations that shape the object in the very act of making it knowable. Collectively, these studies reveal that representations are contingent products of layered work in which instruments, categories, and conventions select what becomes visible, define what counts as the object, and embed evaluative assumptions about what is worth measuring at all.

The second insight extends this constitutive claim. Representing not only shapes how an object is known but also reshapes the object itself. Mol (2000) argues that the practices of measuring blood glucose and rendering a diagnosis *enact* diabetes, influencing what patients eat and how they perceive themselves and thereby reshaping the very condition under investigation. Heuts and Mol (2013) make a parallel observation about what constitutes a "good tomato," showing that different "registers of value" emerged across different audiences and practices, and that informants did not merely evaluate tomatoes against a fixed standard but actively worked to *make* tomatoes good through "sustained and respectful tinkering toward improvement" (p. 125). In each case, the process of making a physical entity knowable simultaneously transforms what that entity is, what it is worth, and what it can become.

The third insight builds on the first two. If representing is constitutive of how the object is known and the object itself, then *how* representing is structured – as a single act or as a sequence of mediations – becomes consequential as well. A growing body of work suggests representing is rarely a discrete event but a layered process in which each stage of mediation

introduces its own ontological assumptions, constraints, and evaluative criteria. Kruse (2021) uses semiotics to describe different levels of digital representation, from raw data through increasingly abstract symbolic layers, with each layer embedding its own choices about what counts as relevant, visible, and actionable. Barley (1996) described the practitioners who managed this process as working at "an empirical interface: a point at which a production system met the vagaries of the material world," operating "with one foot in the material world and the other in a world of representations" (p. 418). The work at this interface is not merely technical but evaluative: quantifying the qualitative allows observers to evaluate objects differently (Muniesa, 2011), and successive rounds of algorithmic manipulation may further transform representations into what Kallinikos (2009) calls a "computational rendition of reality," creating embedded categories, labels, and attributions that enable evaluation and ranking.

Taken together, these three insights paint a picture of representing as a process that is neither neutral nor singular: it unfolds through many distinct steps that shape both what comes to be known as the represented object and how its goodness comes to be evaluated. A critical implication is that the representing process does not converge on a single, unified depiction of a physical entity. Because each layer of mediation introduces its own framings and criteria, the same physical referent can be rendered in multiple, coexisting ways – what we might think of as distinct *renderings* of the same object. Each rendering frames the object as a particular kind of thing: a set of measurements, a trajectory, a position within a distribution, a projected future state. And because each rendering constitutes a different version of the object, each invokes its own evaluative logic and its own standards of worth. The STS literature has shown that different practices of knowing produce different ontological versions of their objects (Mol,

2002), but it has not systematically examined what happens when multiple renderings of the same physical entity coexist within a single representational system, each presenting the object as a different kind of thing and each generating its own pressures for action. In contemporary digital systems, this multiplicity is not an exception but the norm: a single entity can simultaneously appear as a biographical record, an algorithmic prediction, and a comparative position, and the actor engaging with these renderings may find that the object is "good" according to one framing and "bad" according to another.

From *Representation* to *Representing*

The STS literature establishes that the process of representing is constitutive of both representations and the referent. It is laden with ontological assumptions that actively shape what comes to count as an object, what properties are rendered visible, and what evaluative criteria are applied to evaluate the "goodness" of an object. Yet the insights of STS, powerful as they are, leave several questions unexplored that are of direct consequence for organizational scholarship.

First, STS studies have primarily examined how representing constitutes objects within bounded epistemic settings (e.g., laboratories, clinics, scientific expeditions) where the process of representing is typically a discrete event or a circumscribed series of transformations. In contemporary organizations, however, the representing process is continuous, operates at scale, and increasingly involves automated systems that layer multiple stages of mediation in real time. Advances in instrumentation allow measurement of previously invisible processes; increases in computational capacity enable real-time structuring of vast data streams; and predictive algorithms extend the represented object into the future, creating projected entities that have yet to be realized but that influence decision making in the present. Under these

conditions, different renderings of the represented object appear and transform continuously, creating fleeting moments at which evaluation takes place. Where earlier representational systems might have informed a periodic decision, today's systems evaluate and classify physical entities through an ongoing stream of renderings, each producing its own view of the object, each accompanied by its own evaluative assumptions. Christin's (2020) study of digital journalism illustrates this shift: web analytics dashboards do not produce a single representation of an article's worth but a continuous stream of click counts, engagement metrics, bounce rates, and comparative rankings that update by the minute, with editors and writers responding to each rendering as it appears. Brayne (2017) documents a similar dynamic in predictive policing, where individuals are continuously re-rendered as risk scores, hot-spot locations, and network positions, with each rendering triggering its own institutional response. Such settings differ from the bounded epistemic sites of classical STS not in kind but in structural conditions: the representing work is no longer a circumscribed effort to stabilize an inscription but an ongoing organizational infrastructure that generates evaluations and prompts action faster than any single rendering can be deliberated over. Actors may find themselves looking at the same entity but seeing a fundamentally different thing depending on which rendering they engage, and the evaluative question they ask shifts accordingly. How the organizational process of representing – and the multiplicity of renderings it produces – shapes action under these conditions of continuous, multi-layered mediation is a question neither organizations scholars, nor STS scholars have systematically addressed.

Second, while STS has shown that the practices and instruments of representing reshape the physical objects they engage – most notably in Mol's (2002) account of diabetes and Heuts and Mol's (2013) tomatoes – these accounts treat the constitutive work as a relatively

bounded transformation: an instrument or practice acts upon an object and, in doing so, shapes what that object materially is. What this body of work has devoted less attention to is the *recursive feedback* that arises when representational systems are embedded in ongoing organizational routines, where the same physical referent is rendered repeatedly over time and each rendering shaping inputs to the next round of intervention. In settings such as digital twins, remote sensing, predictive monitoring, and real-time performance tracking, the assumptions encoded in each rendering generate evaluative criteria that drive physical interventions. Those interventions, in turn, alter the physical referent that subsequent rounds of representing must capture and the altered referent then produces new renderings under modified conditions. Rendering, in this sense, is the mechanism through which representing reshapes the physical. It is the discrete output at which evaluative criteria become legible to actors, prompt action, and thereby connect representational work to material change. Because different renderings invoke different evaluative logics, the interventions they motivate may differ or even conflict. A rendering that frames an entity as promising may coexist with one that frames the same entity as deficient, and the resolution of these tensions produces physical consequences that are not attributable to any single rendering alone. This recursive dynamic – in which representing produces renderings, renderings shape evaluation, evaluation shapes intervention, and interventions reshape the object that is then re-rendered – has received little systematic attention in either the STS or organizational literatures.

Third, and relatedly, the organizational literature has not examined how the *cumulative* effects of multi-layered representing produce consequences that no single rendering would generate on its own. If each stage of the representing process introduces its own ontological assumptions and evaluative criteria through distinct renderings, then the overall trajectory of

representing may exert a selection pressure on physical referents that is greater than the sum of its parts. A physical entity that represented three different ways is not simply described three times; it is subjected to three distinct evaluative regimes, each of which may prompt interventions that incrementally reshape what the entity is and what it can become. This cumulative dimension is difficult to observe in the bounded settings typically studied by scholars of digital representation, but it becomes visible when the representing process operates continuously over extended periods, as it does in organizations that rely on complex digital representational systems to manage physical phenomena.

These gaps point toward a research agenda that bridges the STS insight that representing is constitutive with the organizational concern for how representational systems shape action and outcomes. We know a great deal about how representations shape knowing and coordinate action once they exist. We know considerably less about how the ongoing organizational process of representing – the sustained, layered, evaluative work of constructing and maintaining digital representations of physical phenomena – shapes action and reshapes the objects being represented. This prompts the question: *How do digital representations come to reshape not only what we know about the physical world, but also the physical world itself?*

Methods

Research Setting

I developed this theoretical account through a comparative field study of data-intensive milking operations across six dairy farms in California, conducted between August 2022 and August 2025. To examine the representing process and its consequences for both representations and the physical world, I needed a setting where the full life cycle of representing could be observed and where its consequences for physical referents could be

traced over time. Dairy farming proved ideal because the entire representing process is encompassed within a single organization: the physical referents (cows), the technologies that convert them into data (robotic milkers, activity collars, lab testing), the computational systems that process data into representations (herd management software, predictive algorithms, composite scoring models), and the users who interpret and act upon those representations (farmers) are all co-located on the farm. A researcher could follow the entire process of representing – from observing how data is extracted from a cow to how farmers engage with different renderings of representational data to how farmers made choices about how to intervene based on what the representation indicated – all in a single day.

Dairy cows are also uniquely productive referents for this investigation because they occupy an intermediate ontological status – neither easily malleable nor fixed. They can be bred for specific traits, treated for illness, trained to conform to behavioral expectations, and culled from the herd in ways that cumulatively reshape bodies, productive capacities, and herd composition. Yet they resist change. Their bodies develop according to biological rhythms, and their behavioral dispositions are not fully programmable. This intermediate status makes the material consequences of representational evaluation concrete and observable while ensuring that the work required to achieve alignment between representations and physical realities remains visible. The representational systems on these farms exhibit a parallel profile – hardware is largely fixed (robotic arms have finite reach, sensor calibration ranges are preset), yet software is partially configurable (farmers can edit data records, adjust algorithmic thresholds, and choose which reports to prioritize). This mix of rigidity and flexibility in both biological and technological systems created conditions under which the tensions and reconciliations between representations and their referents could be empirically traced.

Site Selection and Technologies

I selected six farms from a larger study of twenty-one dairy farms because they employed three data-intensive technologies central to constructing digital representations of cows: genomic testing, activity tracking collars, and automated robotic milkers. These sites operated under similar market, regulatory, economic, and geographic conditions and implemented comparable suites of sensor-based technologies, allowing me to observe how the same representational infrastructure interacted with different farming practices. Genomic testing used proprietary models to predict heritability across more than thirty traits (e.g., milk yield, fertility, disease resistance, and body conformation) aggregated into composite merit indexes that farmers used to stratify breeding and culling decisions. Activity tracking collars continuously monitored movement, rumination, and eating patterns, detecting deviations from individual baselines that signaled illness or ovulation and generating alerts directing farmer intervention. Automated robotic milkers allowed cows to voluntarily enter a milking station where a laser-guided arm located each teat and extracted milk while built-in sensors analyzed milk from each udder quarter for conductivity, flow rate, composition, and somatic cell count, building detailed production and health profiles over time.

These technologies were paired with digital databases, predictive algorithms, and data visualization software platforms to generate composite scores and predictive alerts – such as merit scores, health indexes, heat probability, mastitis likelihood, and sick chance – transforming continuous biological processes into formalized, visible, and actionable metrics. Together, they produced comprehensive digital representations of each cow that organized how farmers came to know, evaluate, and intervene upon individual animals.

Data Collection

The project began as a collaborative, inductive investigation with a broad research question: how do new technologies and data change the practice of farming? For the broader project, I was a part of a large research team that employed a relational ethnographic methodology to observe complex interactions between entities around technology use (Anthony et al., 2023; Bailey et al., 2022). Following Bailey, Barley, and Leonardi's (2021) approach for studying complex, concurrent operations, we deployed teams of two to six researchers simultaneously at each farm for week-long visits that covered the full cycle of routine operations. Multiple researchers simultaneously shadowed different individuals through daily routines, capturing interdependent operations while revealing how different actors engaged with robotic systems and the digital representations they produced. The observation protocol involved detailed running notes on farmers' interactions with technologies, cows, people, and documents. I audio-recorded technical discussions, photographed software interfaces, and algorithmic alerts, and collected physical artifacts including protocols and standard operating procedures. Because much of the expert work I observed was cognitive (e.g., interpreting algorithmic predictions, weighing representational criteria against embodied knowledge) I complemented observation with situated questions and semi-structured interviews covering system implementation, interpretation of algorithmic outputs, and decision-making practices. Research teams convened at the end of each day to compare notes and construct a shared contextual understanding, then expanded running notes into comprehensive narratives integrating actions, conversations, and images.

All audio recordings of interviews and observations were transcribed using otter.ai or by hand. Shorthand field notes taken in the field on pen and paper were expanded into detailed

field notes with photos and excerpts from transcripts inserted directly into the narrative when they occurred. Due to the complexity of farm activities and interactions and the extensive data across various media, one hour of observation typically took around ten hours of developmental work to convert to comprehensive field notes. Overall, the team collected over two hundred hours of observations of ninety-six informants across six farms, conducted seventeen semi-structured with dozens more informal interviews in the process of shadowing, each lasting between thirty minutes and two hours, and compiled a corpus of photos, videos, and other supplemental technical documents and media artifacts adding up to over one thousand files.

Data Analysis

While data was collected by a team of researchers, I conducted data analysis and theorizing for this study independently. The analytical process was iterative and inductive, beginning with open coding to capture a broad range of activities, interactions, and practices across field sites (Glaser & Strauss, 1967). The theoretical direction crystallized around a pervasive linguistic pattern: farmers across all six sites repeatedly invoked the language of "good cow" and "bad cow" as they navigated between digital representations and physical animals. What struck me was not merely that farmers evaluated cows, but that the same term – the "good cow" – seemed to mean different things depending on whether a farmer was assessing a cow's physical compatibility with the robot, reviewing her digital health records, interpreting an algorithmic prediction, or comparing her metrics against the herd average. The concept was at once stable in its ubiquity and unstable in its meaning, shifting as farmers moved across different aspects of the representational system.

This observation prompted me to code systematically for every instance in which farmers made evaluative judgments about individual cows, attending to the criteria invoked, the representational context in which the judgment occurred, and the actions that followed. I then deconstructed the representing process itself, tracing how each encounter with the representational system imposed different evaluative demands and generated different forms of misalignment between digital and physical realities. Through successive rounds of coding and comparison across sites, I identified four distinct representational moments – physical, biographical, algorithmic, and comparative – each of which restructured how farmers evaluated cows and prompted different forms of what I came to call congruence work. I also noticed that farmers frequently described how their cows and herds had physically changed since adopting these technologies, and how they deliberately used representations to direct actions that reshaped both digital records and the animals themselves – a recursive dynamic in which congruence work altered the conditions under which the next round of representing occurred.

To ensure analytical rigor, I conducted negative case analysis to refine my understanding of conditions under which different forms of congruence work emerged, held peer debriefing sessions with colleagues familiar with the field sites, and drew on data from the remaining fifteen farms in the broader study to establish baseline comparisons. This allowed me to identify which aspects of the recursive process were attributable to data-intensive representational systems rather than longstanding features of the dairy industry. While physical modification of cows has always been central to dairy farming, informants across all farms reported that the pace of change greatly accelerated with the introduction of extensive data and predictive models. My analysis culminated in a model of representing as a

recursive process in which four representational renderings continuously restructured evaluations and channeled congruence work in ways that reshaped both the digital and physical realities of the farm.

Findings

The representing process unfolded in two analytically distinct phases. The first operated on the physical cow to make her representable at all. Robotic milkers, designed to extract both milk and data, acted as a physical filter. Only cows whose morphology and disposition allowed them to attach to the device could be represented in the system. Cows that interfaced effectively were deemed "good" while those that did not were judged "unfit" for robotic milking and removed from the herd. This inscription bottleneck preceded any digital rendering, but it determined who entered the representational system in the first place. Once past the bottleneck, cows entered as vast agglomerations of data that could be organized and recombined into *renderings* – partial, distinctly framed versions of the same physical referent that coexisted within the larger representation. Farmers encountered dozens of renderings daily, but three carried particularly strong evaluative assumptions and redefined what it meant to be a worthwhile cow: *biographical* renderings of what she had done, *algorithmic* renderings of what she was predicted to do, and *comparative* renderings of where she ranked relative to her herd. Like the inscription bottleneck, each rendering generated its own criteria for which cows were promoted and which were eliminated. Across these four moments, the farmer was looking at the same cow but seeing a different thing each time: a body that fit or did not fit the robot, a record of past events, a predicted trajectory, or a position in a distribution. Each rendering prompted a different question about whether the cow was good enough. The "good cow" was a common term farmers used for cows with desirable characteristics – high milk

production, health, but also quantifiability and comprehensibility – yet the term was applied with high heterogeneity, and what counted as good in one rendering often conflicted with what counted as good in another.

Each new definition of goodness produced discrepancies between representational ideals and physical realities that farmers worked continuously to reconcile. I call this *congruence work*: deliberate efforts to achieve and maintain fidelity between digital and physical realities. Despite these efforts, representations never reached a finalized, perfectly faithful depiction. The unresolvable tension generated continuous, incremental modifications to both systems, and these modifications accrued unevenly depending on *ontological plasticity* – the perceived malleability of a given system, which oriented congruence work toward the intervention points of least friction. A cow's biology could be shaped through feed, environment, breeding, or treatment, but only within the bounds nature permitted; representational systems could be modified through data edits or threshold adjustments, but only within the bounds set by design choices and embedded assumptions. Farmers therefore directed action toward whichever system was more plastic at that moment – editing the representation to fit the cow or editing the cow to fit the representation. Critically, modifications made within any single rendering did not stay there. Changes to the underlying data fed back into the representational system as a whole, influencing other renderings of the same cow. Each rendering "leaked" the consequences of congruence work back into the corpus, forever altering the represented cow. See **Figure 3.1** for a model of the representing process.

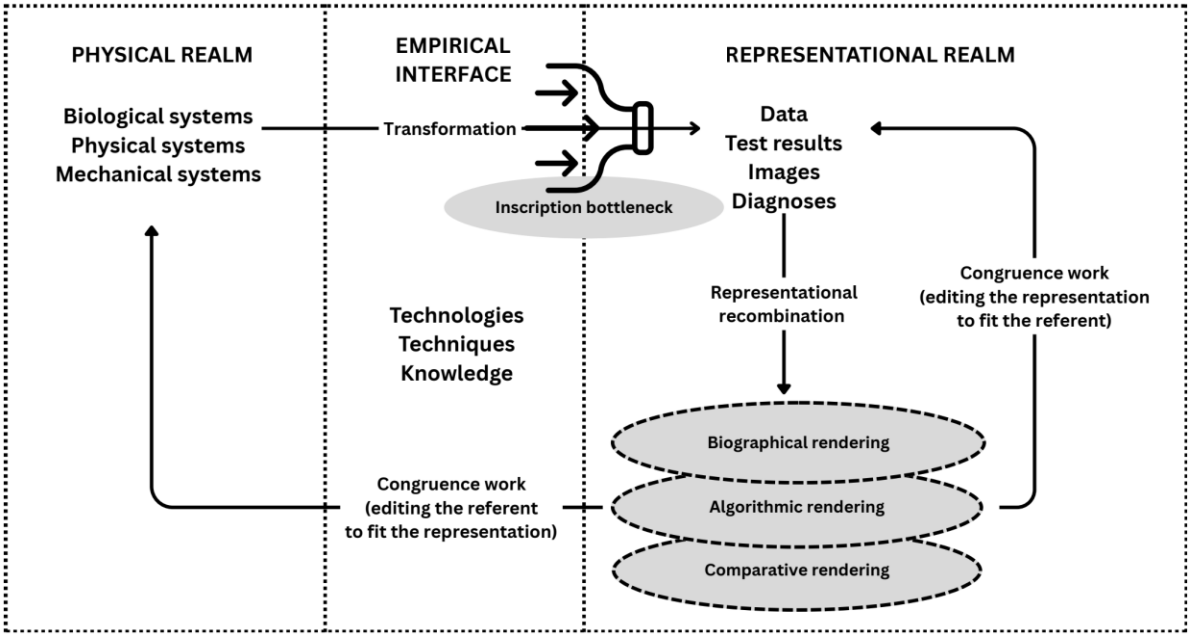


Figure 3.1. Model of the representing process.

In the following sections, I describe each of the four moments during the representing process where the cow becomes a different kind of "thing" to the farmer, the evaluative restructuring associated with each moment, the congruence work farmers deployed to align cows with their representations, and the accumulated consequences of that work over time.

Inscription Bottleneck: Representing requires cow-robot compatibility

Before a representation could exist, farmers first had to make cows legible to the devices responsible for extracting the data on which representations were built. Much like the inscription devices of scientific discovery, the tools used to quantify qualitative information (e.g., robotic milkers, activity tracking collars) imposed physical, behavioral, and morphological constraints that cows had to overcome to become representable in the first place. This phase of the representing process operated directly on the physical referent (not on data) making it the most immediate site where representing began to reshape what a cow had to be.

Rendering a cow into data required two forms of cow-device compatibility. First, farmers had to assess whether each cow would behaviorally "take to the robots" – whether she would voluntarily enter the milking stall and stand still long enough for the robotic arm to attach and produce a clean, uninterrupted data and milk stream. Farmers spent multiple days training cows that had never encountered a robotic milker. Despite these efforts, some cows never adapted. One farm owner of DF10, Fidel, explained this bottleneck effect in the early stages of his transition to robotic milking:

> "There were certain animals that just did not take to the robot. Very, very few, less than 5%. We had to sell those animals. We sold them to other farms, because when we made the robotic decision, we decided that there was not going to be a second option. It was robot or bust, and those cows went to other farms... Thankfully, more than 95% of our cows did just fine in that environment."

The 5% of his herd simply could not adapt to the dramatic difference between conventional and robotic milking and had to be sold because they were either too nervous or too unreliable to be milked consistently. The second compatibility requirement was morphological and posed more rigid constraints. Robotic systems operated on standardized design assumptions about udder position, teat spacing and angle, and body conformation that determined whether milk and data could be extracted effectively. While robots could accommodate most physical variation, certain morphologic features simply did not fit. Carl, another robotic farm owner at DF9, explained just how fickle this constraint could be:

> "It is a robot. When that arm comes under, it maps the cow's [udder] with a laser. Smaller teats are harder for the robot to find, so short ones don't work well. Teats that are like this are tough [crosses fingers]. If they're pointed out, the robot can do it. But ideally, you want them straight down. The biggest issue would be high set rear teats. The robot has a hard time stretching. It has changed some conformation of what we look for."

In conventional milking parlors, human milkers could manually position machines to accommodate cows with atypical morphology by manually adjusting the angle of the milking

cups or holding them in place until they could securely attach. The robots, however, could only adapt so much, and cows in these edge cases had to be culled. Attachment failures jeopardized not only milk extraction but also data extraction. The robot collected data on milk yield, flow, and quality for each teat; inaccurate attachments corrupted these measurements. Chris, a robotic farm manager at DF9, explained how attachment issues could lead to misreadings and faulty data:

> "If you have a cow and it has a little bit of this (crosses fingers), the milkers can [cross the cups]. So, the more you do this [attach the milking cups to the wrong teat], the more you're gonna injure that teat."

Injured teats led to elevated readings of somatic cell count, a key metric used to calculate mastitis likelihood, which could distort the cow's chance of being diagnosed with an infection. Behavioral disruptions similarly produced representational distortions. Cows that kicked off the machines mid-milking interrupted the data stream, generating incomplete or inconsistent records that created discrepancies between representations and physical realities. Bradley, the farm manager at DF1, explained why one faulty milk reading could lead to amplified representational distortions downstream that are difficult to detect:

> "She's kicking off the machine every day, and even our system's not detecting it because her average is always 20 pounds. When she kicks it off within 30 seconds of attaching, her average stays at 20 pounds, even though she has 80 pounds of milk in her... So [the system's] not flagging like, 'hey, go reattach her.' But then when you sort her to beef and you say, 'why am I sending her with only 20 pounds?' And then you mark her and tell somebody to follow her all the way through [the milking process], then you see she milks 80 pounds, and after three days, that affects her new average. So now when she kicks it off, it tells them to put it back on."

Kicking a milking machine off prematurely recorded a low milk yield that set a revised baseline for her milk production. Because milking machines milked based on the time it took to extract the most milk without overmilking (which can cause teat injury), the new baseline set her milking limits for subsequent milkings which prematurely detached. These consecutive low

milk yield readings made her a candidate for culling due to a string of interrupted milkings triggered by one early detachment event. Behavioral compliance was therefore a precondition for representational completeness.

Overall, robotic milkers designed to extract both milk and data from the cows acted as a physical filter, allowing only cows that could be made to physically attach to the device to be represented. Cows with the compatible morphology and compliant disposition needed to interface effectively with the robots were deemed "good" while those who did not possess these traits were judged as "unfit" for robotic milking. This phase preceded any digital rendering of the cow, yet it introduced an inscription bottleneck whereby cows had to either be made compatible with the inscription apparatus so they could be rendered digitally representable or be removed from the pool of referents altogether.

Representability Restructures Evaluations of a "Good Cow"

Failures of physical and behavioral alignment directly undermined the digital representations upon which downstream activities were based. Consequently, farmers had to carefully evaluate the fitness of cows before they entered the representational system which reshaped how farmers defined a "good" cow. The representing process, at its very first stage, was already selecting for particular physical trait, redefining "goodness" around the requirements of inscription rather than production alone. This was most apparent in the vocabularies farmers used. Mike, a herd manager at DF12, used specific aesthetic vocabulary to describe this robot-centered evaluation:

> We walk by a cow and [the herdsman] points to a particular cow – "Bonita" ['pretty' in Spanish], he says. I ask him which cows he thinks are 'bonita.' He points to several individuals with different udder shapes and teat placements. These cows are 'bonita' because they suit the robots well, he explains in Spanish. [excerpt from field notes]

Another farm manager at DF11, Preston, echoed this aesthetic evaluation and specified that once her udder dropped too low due to age and repeated lactations, she would eventually become incompatible with the robots:

> "She's a 'big cow, beautiful cow' in her 4th lactation but they won't breed her again and will likely sell her when her milk production drops because the robot won't be able to milk her once her udder hang too low. She might find a home at another dairy, but she can't stay here."

Beauty – and with it, positive evaluation – became associated with morphological compatibility with the robots.

Certain amenable dispositions also became bases for selection. Cows who repeatedly failed to visit the robots were deemed "lazy" for requiring constant fetching (physical prompting to go to the robotic milkers), as Daniel, a robot barn manager at DF6, noted: "Some cows are just lazy and will pop up on the list every two weeks. Those individuals are just reluctant to move around sometimes." Behavioral evaluations abstracted to breed-level stereotypes as well, as Dr. N, a veterinarian at DF9 explained the breeds of cows most amenable to robotic milking:

> "It's a known thing about Jersey [cows] that they're more curious. I feel like curiosity makes them start going in the robot, start eating grain in there. Some Holsteins are a little bit skittish, so they're not so amenable. Jerseys are also stubborn, unfortunately (laughs). They're very habitual as are all cows but you could argue […] that they're better for robots than Holsteins."

Evaluative restructuring at the inscription bottleneck thus produced a new criterion of goodness: a good cow was one whose body, behavior, and temperament suited the representational apparatus and could be reliably converted into data. Cows that failed this test were removed early in the representational pipeline.

Congruence Work Aligns Bodies and Robots

Farmers had an interpretive role in deciding when to modify cows to fit the robots and when to adjust robots to accommodate cows. Congruence work targeted one or more of these avenues depending on how farmers perceived the plasticity of each system.

Where technology was less plastic, farmers modified cows. Robotic hardware was largely fixed – stall width, arm reach, sensor calibration ranges, and attachment protocols were not easily modified. Farmers largely worked on modifying cows' bodies and behaviors to meet the robots' requirements. They trained reluctant cows, manually fetched them, and experimented with retraining protocols. One example from DF12 illustrates how they approached this behavioral slippage problem:

> He [farm manager] selects from a dropdown list ('1 milking yesterday') to see which cows need attention. This list is slightly different from the fetch list because it shows the trained cows that failed to report to the robots and must be retrained... Mike came up with the idea of retraining them periodically if they started to show slippage in desired behavior rather than transferring them to the conventional milk herd. They are experimenting with this new method, and he says that they have had success with it so far. [excerpt from field notes]

Training cows to visit the robots reliably was a core activity, both for milk production and data accuracy. Yet behavior alone did not determine representational viability. Farmers compensated by physically bringing cows to the robots when they did not go on their own. Daniel from DF6 described some cows that he accommodated despite their failure to comply:

> "[One cow] needs to be fetched every day, but she is a good cow and gives 94 pounds of milk and is 194 days pregnant. 'She just doesn't like the robots,' he tells me, laughing... 'She'll stand right there until I push her in.' He says it's easier to fetch her every day since she is used to the robot than to move her back to the conventional herd and bring a brand new cow in to train."

Here, Daniel decided the cow was "good" despite her delinquent tendencies – she produced a lot of milk, was nearly due to give birth, and had already acclimated to robotic milking. The

only issue was that she needed to be reminded, which he accommodated by physically fetching her.

Where technology was more plastic, farmers modified the device. Farmers could "train" robots to map cows with unusual udder shapes through "supervised milking" – physically intervening during the attachment process to help the laser positioning system find teats more efficiently. This practice was typical across all robotic farms, including DF10 where this practice was a systematic part of their daily routines:

> [The herdsman, Fabricio] walks up to a cow that is near the entrance to the robot and waves his arms to encourage her to enter the robot. Once inside, he starts a supervised milking sequence. The cow seems a bit uneasy, turning her head to look back towards the robot arm as it moves into position under her udder. She keeps moving back and forth as the robotic arm tries to position the milking cups. At one point, she steps her hind feet over the milking arm, so it is now completely behind her instead of under her. The articulating arm moves back and forth a few times to find the correct position while Fabricio bends over next to it and tries to get the cow to stand still. After a few minutes of trying, the robot is unable to attach and the cow is released. He exits the robot room and ushers the cow around the pen back towards the entrance to the same robot to try again. She gets in and he performs the same process by initiating a milking sequence and watches closely. The robotic arm repeats the laser scanning process, and the cow seems a bit less agitated this time around as Fabricio stands at her back end next to the robotic arm and gently strokes her back leg with his hand. He tries to manually position the milking cups by touching the buttons on the touchscreen to move the arm back and forth until it is right under the teats and then deploys the cups. This time, the milking cups attach successfully and the cow seems more comfortable. [excerpt from field notes]

This process required significant involvement on the farmers' end to train the robot while simultaneously managing the cow over multiple days to facilitate better machine learning and cow comfort. But sometimes even this heroic effort was not enough to achieve cooperation between the cow and the robot, in which case the cow ultimately was removed from the herd.

When the cow was too incompatible, representation ended. If training and supervision could not align the cow to the device, the representational process broke down entirely, as explained by Fidel, the farm owner at DF10:

> "The other thing with a robot system is the udder conformity or the teat placement has to be pretty good. It doesn't have to be perfect. The robot's quite incredible for detecting and milking an animal out. But if you have a cow that has back teats that are crossed, the robot isn't going to be able to attach. So, there were a few animals in the herd that we had to sell because of those issues as well."

This was a common experience across all farms during the transition to robotic milking. Because data extraction devices were less malleable, congruence work at this bottleneck manifested mostly as cultivating and selecting for robot-friendly cow characteristics. Only cows whose bodies could be reliably turned into data remained part of the herd.

Consequences of the Inscription Bottleneck

Congruence work at the inscription bottleneck nudged herds towards being more robot-compatible over time, producing unintended consequences for cows and farm operations. As farmers removed cows with incompatible teat conformation and disposition, herds became more morphologically uniform and docile – not only through selective inclusion, but through new farmer-cow interaction patterns that emerged as cows were reconfigured to fit the robots. Daniel, the robot barn manager from DF6 drew explicit distinctions between the cows from the conventional milking operation on the same farm and his robotic herd:

> "[Conventional cows] are more wild. Here [in the robot barn], you can walk by them, pet them, and they don't move. They know I'm not going to hurt them. They just know they need to get up to get milked. Over there, they get shots [vaccines, hormone injections] so they're all jumpy. You can't touch them because people are touching them all the time. They're calmer over here."

Daniel worked with both herds and the stark contrast between the cows was a point he continually brought up. Chris, the farm manager at DF9, echoed the same observation as he

described the difference between his herd before and after they implemented robots (painting an adorable image of his curious cows 'licking you to death', a claim that I personally verified to be accurate):

> "People think that it's just a robot that milks the cow... No, it changes everything. It changes your cows' behavior. Before you would go, and if you stood next to the fence and dropped the chain on it, you'd see all the cows get up because that was normally a sign that, 'Hey, it's time. We got to go with the herd to the barn.' Now you go in there, and they'll probably try to lick you to death! People don't hang out in their pen anymore."

Cows in robotic operations appeared physically and behaviorally more relaxed. Herds also became more morphologically uniform since farmers were removing those with physical features that were less compatible with the robots' immutable design features. Again, many farmers like Carl from DF9 described the material consequences of misfits with the machinery:

> "[Cows that don't fit the robot] have one of two options. They would go to beef [sold to a slaughterhouse], or we sell them to conventional dairies depending on what the conformation problem is. We probably lost about a half a percent [of our herd]in the changeover because conformation wise, they didn't fit."

Furthermore, this selection pressure continued beyond the initial transition to robots. Since young cows were continuously introduced and existing cows' bodies changed as they aged, cows had to consistently satisfy the physical requirements for robotic milking throughout their careers on the farm. The representation process, even before data collection, physically reshaped the population of referents available for representation.

Biographical Rendering: Representing expands what is visible

Once cows passed the inscription bottleneck, they entered the representational system as data that could be organized into renderings. The biographical rendering consolidated raw measurements into digital insights about cow medical history, daily activity, and milk production. In this rendering, the cow appeared as a longitudinal record of what she had done

and what had happened to her and the evaluative question shifted from "can this cow be made into data?" to "what does the data reveal about this cow?"

Prior to data-intensive technologies, farmers relied on embodied observation or rudimentary histories recorded as discrete entries – often only 30 or so events across a cow's lifetime. Data-intensive technologies exponentially expanded these sporadic snapshots into continuous, individualized, and granular streams of insight. Activity tracking collars equipped with accelerometers and microphones rendered movement, resting, rumination, and feeding data into 24/7 time-series graphs traceable across months. Robotic milkers equipped with chemical analyzers tested somatic cell counts three to four times daily per individual cow rather than bulk-tested once daily for entire herds. Processes such as lactation fluctuation, milk conductivity changes, and body temperature became visible in ways physical observation alone could never detect. Fidel, the farm owner at DF10 explained the affordances of this expanded data succinctly:

> "The collars give us an incredible amount of data. For one, they identify the animal when she's in the robot. Two, they track activity. It's kind of like your Apple Watch or Fitbit. And then, because of the position where they're located, they give us the rumination or the eating time. It really lets us know the performance of the animal... I can show you on the computer (pulls up a cow's activity report on screen). We know more about that animal than we know about each other. It's really, really cool!"

Digital traces made subtle changes perceptible, interpretable, and actionable. A cow that appeared fine in the barn might show a steady upward trend in somatic cell count (a subclinical sign of potentially deadly mastitis) allowing farmers to intervene before symptoms emerged. As Carl from DF9 explained: "If the robot says a cow has mastitis or she's sick, the robot knows what it's talking about. It's catching animals that are probably subclinical a day or two earlier than we visually catch them." With new visibility came new diagnostic and evaluative

baselines that reshaped not just what farmers knew, but what they perceived as relevant and worth attending to.

Clean, Comprehensive Records Restructure Evaluations of a "Good Cow"

The biographical rendering redefined "goodness" as a function of longitudinal and comprehensive historical insights. The contrast was visible when comparing how farmers with sparse versus comprehensive data evaluated cows. With limited biographical data, evaluations were blunt. Damien, the hospital manager at DF6 demonstrated this evaluation process while reviewing candidates for culling on DairyComp, the farm's herd management system and central data repository for everything cow related:

> Farm #1 (conventional): Damien approaches the main computer and pulls up the DairyComp database. He shows me how to look up specific cows and their records. He shows me one cow with a long list of entries in her record. "One here, two here, and three right now…She's back again. Not good cow," he says as he counts the number of positive mastitis infections requiring antibiotics and hospital stays this particular cow has accrued since November. He points at the screen and tells me her milk production is low (WMLK1 = 65; the cow above had a WMLK1 = 109). "For me, beef," he says, confirming that this cow is not suitable to stay in the herd and will be sent to slaughter. He pulls up another example. "Look at this cow. She's 190 [days] in milk. She's open, DNB [Do Not Breed]. She's not good cow, not very good milk. That's automatic beef." [excerpt from field notes]

Damien utilized limited biographical data, such as days in milk, weekly milk average, last milk yield, trips to the hospital, to evaluate whether these cows were "good" cows or "not good" cows. With more comprehensive biographical data, evaluations became more nuanced. Daniel, the robot barn manager at DF6 had more extensive data available to him through activity tracking collars and robotic milkers, which he reviewed for each mastitis and heat prediction:

> Farm #2 (data-intensive): "This [activity] graph helps. It gives you more data to decide [which cows are actually in heat]. You see the activity of cow #49210? It went up and it just came back down slowly compared to the activity of #54380 which just went up and back down right away. That one (#54380) is not in heat because she just went up and then right back down. This line should go up fast

> and then come back down slowly. But this (points to the green activity line for cow #54380) went back down really quick. She's not [ready to be bred]. [Similar with assessing mastitis], She's ten and ten (high mastitis likelihood score), but her milk is right on (points to the numbers in a data table that show average milk yield and the previous day's milk yield which are not that different). Nothing's wrong with her."

Discrete data entries about mastitis events, milk production, and reproductive status were still important benchmarks for assessing the "goodness" of a cow, but available biographical information was expanded in both depth and breadth with increased data capture and granularity, allowing farmers to evaluate and compare trends over time and recognize patterns that informed more nuanced decision making. This level of insight meant that a "good" cow was one that appeared productive and unproblematic across her accumulated record, as Preston, the farm manager at DF11, explained:

> "A good cow is a cow that gives a lot of milk and doesn't give any problems…. Like problems with the feet, problems with mastitis…. Let me see if I can pick out one of these cows (turns to Horizon to search for a cow and types in an ID number from memory – pulls up cow #1025). So, this cow is a good cow. He clicks to her "Events" tab (photo 35) and shows me that she has had two breedings on the ov-sync program and has lost one pregnancy. But other than that, she has had no cases of mastitis even into her 5th lactation. "She's a big cow, a beautiful cow," he tells me. He pulls up her milk quality report (photo 36) and shows me her milking stats. Her production is approaching 140 lbs. of milk and has been steadily increasing over the last few weeks."

Preston and other farmers typically reviewed as much of the available data as needed to make the most informed decisions. Given the amount of data being collected with these data-intensive technologies, if a cow's biographical rendering did not show common indicators of being a "problem" cow, then she was deemed "good." Goodness was redefined as the absence of red flags across a comprehensive longitudinal record – a standard that only became possible through the expanded visibility that the representing process itself produced.

Congruence Work Aligns Embodied Observations and Digital Insights

As new data insights came into view, discrepancies between digital and physical realities surfaced. Farmers chose to either adapt the rendering to reflect the physical cow or manipulate the cow to fit what the rendering indicated.

When biographical rendering was less plastic, farmers modified the cows. Farmers often encountered situations where the rendering told one story and the physical cow told another. Damien, the hospital manager from DF6, described this ambiguity that arose when a cow's physical appearance did not match what he saw in her digital record: "Sometimes, you think 'Oh, she's a good cow'," he says, pointing at the screen, "but then you go check the cow and [she's] skinny and looks like this," he explains while miming a sickly-looking face." In these cases where cows and renderings were not aligned, farmers had to find a way to make them align. When the representation was rigid or indisputable, congruence work meant editing cows based on their representation – culling those with digitally documented reproductive failure, chronic illness, or persistent deviations. Preston, the farm manager at DF11, explained how this practice shaped his breeding and culling decisions:

> "We won't breed no more because we bred her a lot of times already, spent a lot of money, and the cow's not getting pregnant. When the cow says, 'keep open,' we don't do nothing else to breed it. We leave the cow to give milk and after she goes down, we sell it. You can see this cow's days in milk, 450 days… Last week, we sold a cow that was 800 days in milk. She was a really good cow because she gave all this milk, but she was not getting pregnant, so that is why we sold this cow."

These decisions were triggered by data-backed insights that made cumulative problems visible.

When biographical rendering was more plastic, farmers modified the data. Sometimes, congruence work looked like editing data when physical observations contradicted the data if the cow could not be modified. In one instance, Dr. N, the veterinarian for DF9, who

was performing a pregnancy confirmation via ultrasound determined that the fetus' size did not align with the insemination date in her digital record. This conversation between the herdsman assisting her with vet check and herself revealed how they reconciled these discrepancies:

Veterinarian:	Is she a preg 2 [2nd pregnancy check] or previous [1st pregnancy check]?
Herdsman:	Oh, she's still a previous? (visibly surprised at the question)
Veterinarian:	Yeah…
Herdsman:	Give me a number. Let's see how good you are.
Veterinarian:	Let's go with 48 [days].
Herdsman:	(Looks at records on device) Nope. I can give you 52. (laughs)
Veterinarian:	Really!?
Herdsman:	Yeah, 33 or 52 [days pregnant]. You choose.
Veterinarian:	Ok, let's keep it at 52 [days].
Herdsman:	(Turns to me and shows me the device) So, I can 'allow' [this change]. (he selects the older breeding date from the two options given in the app)…Okay, so now [Dr. N.] comes and does the ultrasound and tells me that right date. So now I can change this.
Veterinarian:	We just select 'preg from previous date' and it gives him options.
Herdsman:	It gives me numbers of what it could be. So, I just change it to whatever she told me, and I just accept it and she's [recorded as] pregnant [from this earlier date].

Because the physical reality was fixed (i.e., they could not change the fetus's age) they adjusted the insemination date, but *only between the two date choices the system provided*. This demonstrates how plasticity allowed farmers to modify the rendering, but still within the constraints of the system.

Consequences of Biographical Rendering

Congruence work in biographical moments shaped both what data was incorporated into representational systems and what physical interventions followed. Some farmers used biographical data in aggregate to identify trends that prompted new practices, as Carl from DF9 explained:

> "We breed them between 18 and 24 days, that's where 53% of our breedings are. Our conception rate goes up to 43% with 95% confidence intervals – pretty tight there. We're looking to see if [the breeder] is breeding the cows [within this window]. When we were pretty aggressive on breeding before robots, we got up to 14% out of cycle. That was a big focus with the breeders to check dates. You've got to make sure you really think she's in heat before you breed her because you're breeding out of cycle and that's hurting your overall conception rates."

Insights like this often led to new congruence work practices and physical interventions (in this case, checking each cow's biographical data to make sure she was within the expected window for ovulation). Biographical insights shaped which cows farmers chose to keep and which they chose to sell, as Bradley at DF1 described:

> "So, you can go through and look at these cows one by one… You look her up – say cow 5941 – and without me even seeing the cow, I can look at the numbers and say, 'yeah, let's send her to beef.' Especially when you dig into her history. She's only had mastitis one time, she's not chronic – but then there are other cows where you click on them and she's had mastitis over and over. She's a 56-pound cow. Why do I want her back?"

If a cow's biographical rendering looked undesirable, she was removed, slowly shifting herd composition toward individuals with "clean" records. At farms with detailed digital dossiers accessible in real-time, decisions could be made more precisely and expeditiously, shaping herds in a more directed manner over time. The representing process, through biographical rendering, expanded insights by which cows were evaluated and, in doing so, reshaped which cows remained in the system to be represented downstream.

Algorithmic Rendering: Representing Imbues Data with Meaning

Algorithmic rendering presented versions of a cow that did not necessarily exist yet. Machine learning models and predictive algorithms took biographical data and converted it into predictions, risk probabilities, composite indices, and categorical labels that assigned meaning to metrics. In this rendering, the cow appeared not as a record of the past but as a

projection of the future and the evaluative question shifted from "what does the data reveal?" to "what does the data say about what this cow will become?"

By transforming metrics into higher-order interpretations (e.g., sick chance scores, mastitis risk probabilities, heat likelihood, genomics indices, and composite productivity ratings), algorithms provided predictions about a cow's probable future and recommendations about how to act. Farmers encountered dozens of these algorithmic renderings daily through reports such as "Breed and Heat" lists predicting ovulation likelihood, "Cull cows" reports recommending removal based on projected performance, "Udder Health" reports with color-coded probability scores, and "Attention Cows" lists flagging animals for review. All came embedded with symbolic boundaries delineating cows into groups or categories.

The most striking example was genomic scores – products of generations of genetic and phenotype modelling that used DNA analysis to predict an animal's genetic potential for production, health, and fertility early in life. Proprietary algorithms provided indexes like the "Dairy Wellness Profit Index" to rank animals on overall profitability and health as predicted for their lifetime, including susceptibility to disease, milk production potential, and even attributes about body shape and constitution. Fidel, the farm owner at DF10, explained how he used this data in his breeding and reproduction strategy for his farm:

> "So let's say that we genomic test an animal, and maybe she doesn't have all of the genetic potential or performance that we want, then we'll breed that animal with that Wagyu beef, so that we don't receive an offspring from that animal, and then the best animals in our herd will get sex posting or Jersey semen, so that we do get a female that will come back into the herd that has those genetic traits that we're looking for."

This level of artificial selection at the level of genetic traits created new evaluation criteria that never existed before the introduction of DNA analysis and predictive genomic models.

Promising Futures Restructure Evaluations of a "Good Cow"

The evaluative criteria embedded in algorithmic renderings reshaped how farmers thought about "goodness" – even down to a genetic level. The "good cow" was no longer merely one who performed well today; she was one whose modeled trajectory aligned with thresholds, efficiency targets, and probabilistic futures that farmers deemed desirable. Future potential became the measuring stick by which farmers like Carl evaluated the goodness of their cows. In this case, he described a predictive model that calculated the likelihood that a cow would become ill within the next few days based on data from activity tracking collars and robotic milkers:

> "This score tells you the cow's percentage 'chance' of being sick. So, it's using activity, rumination, time eating, minutes laying down, but it's also looking at the milk quality stats. That's an algorithm built into the Horizon system."

These predictions allowed farmers to prospectively intervene with supportive care or treatments. With genomic scores, farmers predicted a cow's potential to become a "problem cow" by assessing genetic probabilities. Daniel, the robot barn manager at DF6 explained how he incorporated genomic scores into his decision making around which cows to introduce into the robotic herd:

> "I just recommend we go by genetics. The higher the number, the better the animal will be during her time here [at the robot barn]. This cow is gonna be a good cow… 737 – Her genomics are really good. The higher the genetics, the less she'll get sick, the less she'll get lame, the less she'll get mastitis. You want to bring those [into the robot barn]. But if you can't find the higher ones, you start bringing in the lower ones – the 300s.…. We make sure all the best cows are over here… We want the best cows here so they can do their own thing. We breed them all to sex semen to get better milk cows out of them."

For Daniel, he only wanted to bring in cows with the best predictions for milk production, reproductive success, and vigor so she could achieve her full productive potential under robotic milking. It is important to note here that *these cows had never even produced milk*. He was

evaluating heifers (juveniles that have not reached milk productive age) based on genomic models. These models predicted a cow's future performance by reading her DNA at thousands of points and comparing the resulting genetic profile to those of cows whose lifetime outcomes – milk yield, fertility, disease resistance, longevity – were already known. These cows were being evaluated on predictions of future performance that had not yet materialized.

This shift in evaluation when dealing with algorithmic renderings signaled that models acquired interpretive authority. A cow designated as high-risk might be monitored more intensively, treated preemptively, or removed earlier than she would have been otherwise. A younger cow with promising genetic scores might be retained despite mediocre current performance, while the genetic line of a cow with a strong production history could be terminated based on low genomics scores. The object of evaluation had shifted from present condition to modeled trajectory, and "goodness" had become a forward-looking, probabilistic designation, rather than a description of who she was today.

Congruence Work Aligns Present Realities and Predicted Futures

Algorithmic renderings generated misalignments that required active congruence work. Unlike biographical events, algorithmic systems were often less transparent and less malleable, thresholds and modelled predictions were preconfigured by vendors or tied to industry standards, encoding assumptions about acceptable risk tolerance, efficiency, and productivity that were not always visible to farmers. A cow could appear physically healthy yet be algorithmically classified as high risk or low potential.

When algorithmic rendering was less plastic, farmers modified the cows. When farmers could not adjust thresholds, they were forced to either ignore predictions or adjust

cows to fit representations. When ignoring predictions, farmers like Daniel drew on embodied and sensory judgment to make decisions while bypassing the algorithm:

> "We were finding cows that were in heat before the computer could find them. I called the lady [customer service] and talked to her about it and see if we could change it [the threshold]. She explained that's the way they do it based on a study that they did that determined the best time to breed the cow so they couldn't change it."

In this case, Daniel chose to manually separate cows for breeding that demonstrated signs of heat but were not flagged by the system, thus bridging the gap between the representational system and observed realities in the barn.

Alternatively, when algorithmic plasticity was low and farmers' confidence in their own judgment was outweighed by confidence in the system, cows were more frequently acted upon to cohere with representations, seen in instances where farmers blindly deferred to algorithmic predictions despite them being incorrect. Cameron, a herdsman at DF9 noted these mistakes during routine pregnancy checks: "This cow was pregnant. So, the breeder came and he thought she was in heat. The robot held her. It thought that she was in heat, but she was already pregnant but he bred her anyways." He noted that some cows were already pregnant when the breeder inseminated them again based on algorithmic recommendations. While not harmful to the cow or fetus, mistimed or erroneous repeat breeding was costly and unnecessary.

When algorithmic rendering was more plastic, farmers modified thresholds. Many of the modeled outputs were inscrutable, abstract, and fixed. However, some farmers discovered they could change certain thresholds to capture the cows they wanted under a certain label. For example, Marco, the farm owner at DF12 realized he could adjust what counted as a "sick" cow by modifying algorithmic parameters in the software program:

> "I would go over to DairyComp to look at her record to see what's going on with her and it would say 'ILL, ILL, ILL, ILL…' and I'm wondering why she never sorted. Then I'd look at her milk production and see that it was no good.

> It's gone down, it was good, but it's gone down because she's sick. [But the robot] wasn't sorting her so that's how we figured out that we need to increase that sensitivity. They [DeLaval] were a little hesitant to increase it because they said we were already on the sensitive side. But I'm like, ' but the data's telling me she's sick!' so they ended up doing it and I think we're kind of where we need to be now [in terms of the threshold]."

Notably, events like breeding, mastitis diagnoses, and milk production that Marco noticed in the cows' biographical renderings hinted at an issue with the sensitivity of the algorithms that predicted illness, leading him to contact the vendor to decrease the threshold of inclusion. In turn, biographical data were also altered by changes made to algorithmic thresholds that alerted farmers to heat, sickness, and mastitis by directing farmers to breed, treat, and cull according to algorithmic recommendations. In essence, the influence of algorithmic rendering on evaluation was often folded into the representational corpus only to emerge in other renderings downstream. This effect was demonstrated most clearly in one instance where Derek, the farm owner at DF6, lowered the threshold for detecting and diagnosing early pregnancy via a blood test that measured proteins released during pregnancy. However, he made this change without notifying his veterinarians, Dr. E and Dr. S, who noticed an unusually high rate of "abortions" in the herd at pregnancy check:

> This 'grey area,' [the vet] tells me, can be manipulated depending on where the threshold for pregnancy is set. She [Dr. E] says that [Derek] has changed this threshold in the past in [the system], which resulted in many more 'abortions' being noted at P2 preg check, but these cows were likely never pregnant to begin with because they were miscategorized as pregnant due to their borderline blood results. "We get more P2 aborts when [Derek] decides to mess with the threshold." [excerpt from field notes]

In this case, Derek's changed the threshold and failed to update the vets who found out by accident when doing pregnancy checks. They noticed that they were seeing a higher-than-normal number of "abortions" than what was expected for this population at this stage. When they told the owner, he confessed to having lowered the threshold so the system would include

more borderline cases and, hence, more false positives. Not only did the threshold adjustment change the reproductive status of the cows within the representational system (i.e., pregnant vs not pregnant), but it also induced potentially erroneous abortion diagnoses by the veterinarian weeks later.

This had two significant impacts on the cow's "goodness." First, cows that were not pregnant were mislabeled as pregnant, and second, cows that may have originally been falsely diagnosed as pregnant became "aborts" later. This had major downstream consequences for culling and breeding decisions because abortions were considered an indicator that a cow was not reproductively efficient and, therefore, not "good" to keep in the herd or continue breeding for replacement milk calves. These blemishes on a cow's record could have them sold down the line or have their genetic lines terminated. Similarly, cows that were algorithmically flagged as having mastitis would often be diagnosed and proactively treated, which would be recorded as medical interventions that shaped future evaluations and culling decisions.

Consequences of Algorithmic Rendering

Over time, these patterns of manipulating thresholds and editing cows shaped herd composition in ways that extended beyond any individual decision. Breeding strategies shifted toward traits associated with favorable algorithmic projections. Animals flagged repeatedly as high-risk were removed earlier. Thresholds were adjusted to construct the representational ideals that farmers envisioned while simultaneously directing choices to conform to representations in a recursive manner. Farmers explicitly recognized this recursive dynamic, noting how evaluation standards rose as herd baselines improved, as Carl exclaimed:

> "The cutoff point for [inclusion in our milk herd] was down here about three years ago (holds hand up near his face). Now, it's here and continues to rise (holding hand up above his head). Those animals that are below the cutoff are sold to other herds in Texas… [One year ago] we were at 55.9 pounds of milk

> [per cow per day], 2.6 milkings, 4.91 fat, 3.55 protein... [Now,] we're a year later. Have we made genetic progress? Have we made facilities [changes]? Are we milking the right cows? We were at 55.9, we're now at 59. Those cows are up on fat. And we're going up on protein. So yes, we're still milking the right cows. We've increased the milking. So, the guys are doing a good job of getting the cows trained to go to the robot... The genetics have improved."

These "right" cows then acted as the physical referents for the next generation of digital representations, serving as the previous year's data for comparison – and thus the cycle continued. The representing process, at this stage, was not merely predicting futures but actively selecting for them, as the physical entities that survived algorithmic evaluation became the baseline against which subsequent renderings were produced.

Comparative Rendering: Representing Makes Individuals Commensurable

Whereas other renderings presented different versions of the same cow, comparative rendering placed each cow into explicit rankings within her herd, transforming an individual into a component of a population. In this rendering, the cow appeared as a position relative to other cows and the evaluative question shifted from qualities of the individual to "how does this cow compare to others?"

Prior to computational sorting capabilities, farmers assessed cows largely in isolation. A cow producing 90 pounds of milk was considered productive. Farmers understood relative performance tacitly, but cognitive bandwidth limited distributional comparison across hundreds or thousands of animals. Herd management programs computed composite relative value scores, rank-ordered lists, and percentile comparisons that situated each cow within a herd distribution. These relative value indicators became heuristics for farm owners like Ben from DF1 to make simple rules to guide decision making about which sick cows to sell and which ones to treat:

> "DHI has a formula where it calculates relative value. The average cow is 100 relative value. Some cows might be 120. Well, you want to keep that cow around because she's a good cow. When it's time to sell cows for beef or whatever, if there's one that's 84 or something like that, it puts it all into a number. I don't understand how they do it, but you're getting rid of that 84 relative value cow versus the 120. You might spend more money trying to save her or whatever."

Cows ranked below the herd average of 100 were deemed less worth treating while those ranked above average would yield higher return on investment. Interestingly, these rankings were used to make consequential decisions about selling and treating cows, but farmers rarely understood how these scores were actually calculated.

Comparison Restructures Evaluations of a "Good Cow"

Comparative rendering transformed evaluation from absolute assessment to positional ranking. A cow's "goodness" was no longer individually assessed but relationally defined by how she compared with her peers. While the technical details of how these scores were calculated were often opaque, some farmers understood how certain variables influenced the numbers, as Bradley from DF1 explained:

> "Every milk test updates that relative value. There might be a cow that started off at like an 82 relative value, and she's producing 130 pounds. And that's the way she performed for the previous 30 days. Then when she gets tested again, it updates her number and says, yeah, this cow went from 82 relative value, she's up at 105 now. Next test, she keeps performing the way she is, now she's up at 130."

These rankings updated continuously. As a result, each comparative rendering produced a new evaluative frame as new data fed into the system. Under this new evaluative structure, a cow who would have been retained under prior evaluation conditions became expendable by virtue of her position in a high-performing herd as David, the assistant manager at DF6 noted:

> "A good cow is different on every dairy, depending on what the herd average is. In instances of lameness, mastitis, or sickness, there's a number where, if

> she's in another herd, she'd be considered a "bad cow' because she's been lame twice or had mastitis three times."

What would be considered a bad cow at one farm could, comparably be considered a good cow at another. Bradley at DF1 echoed this: "I'm even looking at 70 to 75-pound cows to get rid of, which to someone else might be their top cow."

Comparison introduced competition into evaluation, making every cow's standing subject to revision whenever the distribution shifted. It also changed what made a cow valuable in the first place. Because comparative renderings drew farmer attention to the tails of the distribution – the bottom performers flagged for culling and the top performers flagged for breeding or special treatment – the cows who attracted no attention were those who sat unremarkably near the mean. Those cows required no intervention and never surfaced on a flagged list. Out of this dynamic, farmers converged around a striking ideal of the "invisible" cow, who was "good" precisely because farmers never knew she existed. Bradley at DF1 described this cow as desirably unremarkable:

> "A good cow to me is always gonna be kind of an *invisible cow*, if you know what I mean. Like, you don't know about her because you haven't seen her – because you never had to treat her or mess with her... She'll produce 42,000 pounds of milk in 305 days. If you do the math, that's about 120 pounds a day. She never went to the hospital. So, every single pound of milk went into the tank and I sold it for human consumption... I call her an invisible cow because, honestly, I forgot she was here. She just showed up. Nothing happened. She never required anything."

Daniel at DF6 articulated the same logic. A good cow was one who blended into the herd rather than standing out:

> "I make a list of how good the cow is. How much milk they give, how good they're gonna be genetics and all that. I go out there and look at them to make sure their udders are nice and square... Then you don't even see them unless they pop up on your list for shots or something."

The invisible cow embodied a particular logic that comparative rendering had made dominant: *goodness as the absence of deviance*. She was not the highest producer in the herd, nor the most genetically promising – those cows would have appeared on a different list and drawn a different kind of attention. She was good because the comparative system had no reason to surface her. Yet even cows who had passed all prior evaluative stages could lose this invisibility if the herd around her improved or if she experienced one bad day. Bradley at DF1 described a case when one of his best cows was flagged to be culled due to decreased milk yield:

> "This high relative value cow – she's produced 150 pounds every day, she got bred one time, she became pregnant, she's never been to the hospital. That's like your all-star cow. So, when you're gonna send a cow to beef and her relative value is like 150, you say, wait, why is she in the beef pen? Oh, she's low in milk – well, she might be sick. We need to get her back. We can't sell her. That's a good cow."

After realizing she had a high relative value score, he decided to keep her because she ranked highly against the rest of the herd. Evaluations of goodness were based on cows' positions within an ever-shifting distribution and became a property that depended as much on remaining unseen as on performing well.

Congruence Work Aligns Herd Composition and Representational Goals

Congruence work in comparative moments was directed toward improving herd averages. Because rankings were based on aggregations of individual-level data, comparative metrics could not be directly modified within the representational system. Farmers instead manipulated aggregate numbers by modifying herd composition primarily by proliferating overachievers and terminating underperformers – a form of congruence work that was unique in that modifying the herd simultaneously modified the comparative rendering. Bradley from DF1 explained the logic behind this:

> "You're only gonna be as good as your worst cow, right? So, we used to keep 50-pound cows because we could [cover our costs]. Now we're raising it to 70 [pounds of milk per cow per day]... Hopefully in two years, that number will be 90. And if that happens, our average will be closer to 100, maybe 102. So, this tool [algorithmic system] helps. It creates a list of 40 cows that I want to send to beef, and I get my guys out there to verify [that they can be sold]…. Anything above 120 relative value gets sexed semen. Anything below 90 relative value, we breed to Angus. So, if we have an Angus calf, we sell it on day one. That genetic line is gone. In a way, you're getting rid of your lower genetic animals, and the genetic line keeps improving."

Farmers like Bradley pursued congruence by using distributional data to decide which cows to breed with high-quality sexed semen to replenish their milk herds and which cows they chose to terminate and remove from the herd. For instance, any cow that performed above 120% of the herd was bred to replenish the milk herd, while cows performing at 90% or below would be bred for beef calves, functionally terminating their genetic lineage. Genomic scores were similarly rank ordered, as David at DF6 described a similar heuristic they used to determine cutoffs for breeding:

> "We use the genetic score to choose what cows we're breeding to sex, so we want to duplicate our best genetics and kind of phase out our lower genetics. So, we use that genetic score to pick what cows to breed to sex. So, we'll breed the top 40% of our herd to sex semen, because we want those genetics in our herd. We want to duplicate those cows or improve on those. And then the lower 50-60%. They're still here, we're milking them, we want to keep milking them because they're already here. But we just don't want any offspring from them. So, we'll breed them to an Angus bull, because if they get pregnant, we're still milking them, we're just not putting their genetics back into our replacements."

Ranking cows and selecting them in this way based on simple "if/than" rules allowed farmers to simplify breeding decisions to favor animals occupying top percentiles while preserving cognitive energy.

Consequences of Comparative Rendering

Over time, herd averages gradually rose through these incremental improvement strategies, and the threshold for "good enough" rose with them. Genetic trajectories were

modified based on relative value, with farmers combining comparative insights with other renderings to make breeding decisions that shaped the physical herd. Bradley at DF1 described this cumulative selection effect concretely:

> "Anything above 120 relative value gets sexed semen. Anything below 90 relative value, we breed to Angus. So, if we have an Angus calf, we sell it on day one. That genetic line is gone. In a way, you're getting rid of your lower genetic animals, and the genetic line keeps improving. So [the breeder] has access to that. He'll come in and if she's low relative value, he assigns Angus. If she's higher relative value – and maybe she's had mastitis incidents – he'll assign a bull that maybe isn't top for milk but is really good for mastitis resistance. So that offspring gets the high milk genetics from the mom and the low mastitis genetics from the dad, and you end up with a better genetic animal."

Over generations of selective breeding, the cows produced through these metric-based breeding practices were genetically and phenotypically different from their predecessors. **Table 3.1** summarizes the evaluative restructuring and redefinition of "goodness" for each of the representational moments described above.

Table 3.1. Overview of evaluative restructuring of a "good cow."

Representational moments	Evaluative question	Redefining a "good" cow	Emergent evaluations
Inscription bottleneck	How is data collected? Can this cow be made to be quantified?	Physically fits the robot, dispositionally agreeable and trainable, diligent, and reliably visits the robot, calm and obedient	"Pretty cow" "Lazy cow"
Biographical rendering	What data is available? What does the data reveal about this cow?	Performance data about milk production and quality are in line with desired goals; discrete health and breeding events indicate low maintenance and high productivity	"Good cow" "Bad cow" "Problem cow"
Algorithmic rendering	What does the data mean? What does the data say about the cow's future?	Algorithmic predictions, inclusion/exclusion thresholds, modelled categories, composite genomic scores – all indicate a productive and desirable future trajectory	"Hot cow" "Sick cow" "Genetically good cow"

Comparative rendering	How is the data presented? What relative value does this cow have?	Individual representations situated in relation to herd averages, relative value scores, rank-ordered lists, performance percentiles – all provide new scopes for evaluation	"All-star cow" "Invisible cow"

Summarizing the Representing Process

Across the four moments of the representing process, a consistent pattern emerged: each moment restructured what counted as a "good cow," and each restructuring prompted congruence work that resolved tensions between representational ideals and physical realities. At the inscription bottleneck, goodness meant physical compatibility with the data collection apparatus, and farmers reconciled this through training, culling, or modifying cows to fit the robot – or, when possible, adjusting the robot to accommodate the cow. In the biographical rendering, goodness meant having a data profile that matched desired criteria, and congruence work moved in two directions: editing data when the cow was deemed accurate and editing the cow when the data was. In the algorithmic rendering, goodness meant a future trajectory aligned with embedded predictive standards, and farmers culled, treated, and bred cows based on those predictions while also manipulating thresholds to optimize projected outcomes. In the comparative rendering, goodness became relational (e.g., a position within the herd distribution) and cows were edited based on where they fell relative to aggregate benchmarks that themselves shifted as the herd changed. No single moment produced a stable definition of the good cow. Instead, each rendering invoked its own evaluative logic, prompted its own form of congruence work, and propagated its consequences back into the representational system through changes to the underlying data. The cumulative effect was a recursive process in which evaluations and interventions accrued over time, continuously reshaping both the digital

representation and the physical herd, with neither ever arriving at a final form. This pattern is depicted in **Figure 3.2**.

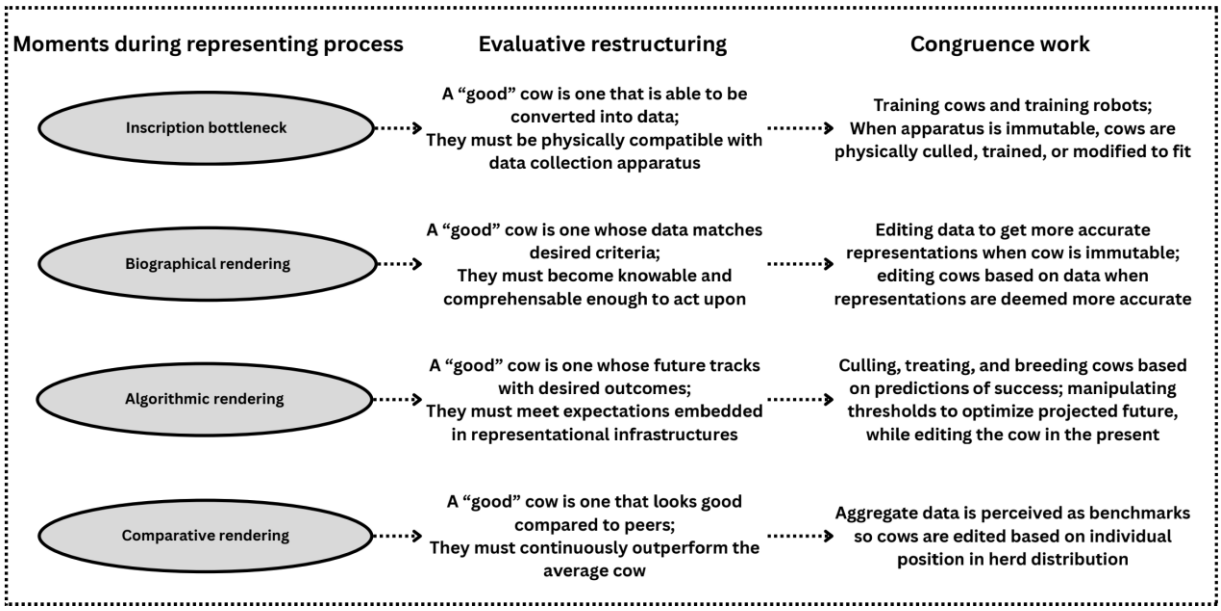


Figure 3.2. Moments in the representing process that invoked evaluative restructuring and congruence work.

Importantly, the recursive process described above was not the product of farmers engaging each rendering in isolation. Farmers routinely combined insights across biographical, algorithmic, and comparative renderings to form synthesized evaluations that no single rendering could produce – considering relative value alongside mastitis history and genomic predictions to arrive at multi-faceted definitions of the "good cow." The representing process thus exerted a cumulative selection pressure that continuously redefined what cows were, what they were worth, and what they could become – not through any single evaluative moment, but through the recursive interplay of multiple renderings, each leaking consequences into the others and each prompting congruence work that physically reshaped the very entities being represented.

The cumulative effect of this process was a form of *ontological co-evolution* (see **Figure 3.3**): two entangled systems reciprocally transforming each other over time. Representations and the cows they depicted reshaped one another continuously, but unlike biological co-evolution, the mediating force was not natural selection but the farmer, whose congruence work repeatedly resolved representational friction at whichever site was most plastic. Over successive cycles, this work reshaped cows' physical bodies (milk production, constitution, morphology), herd composition (genetics, farm-level productivity), and organizational practices (culling and breeding decisions, husbandry routines), producing an ever-moving target in which a cow who qualified as "good" today might become "bad" tomorrow simply because the representational baseline had shifted beneath her. The ontological category of the "good cow" was, in this sense, never fixed – it was continuously reconstructed through the recursive interplay of representation, evaluation, and congruence-oriented intervention.

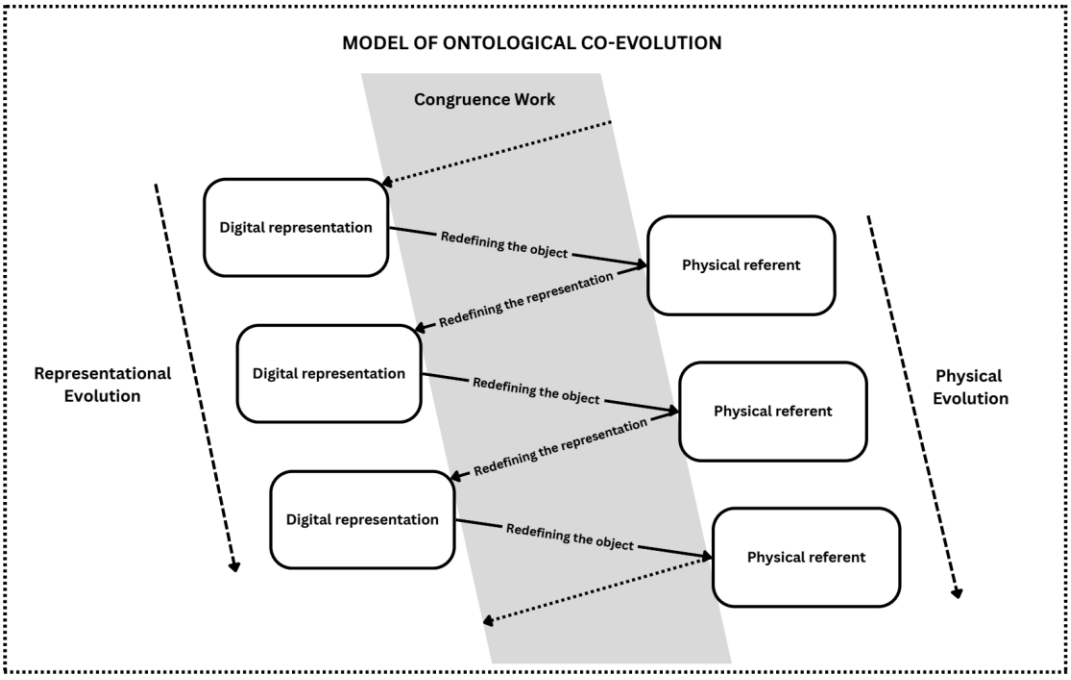


Figure 3.3. Model of ontological co-evolution over time facilitated by congruence work.

Discussion

This study examined how the organizational process of representing reshapes both the digital artifacts and their physical referents over time. Across four moments in the representing process – the inscription bottleneck, biographical rendering, algorithmic rendering, and comparative rendering – I observed how each produced a distinct version of the same cow, invoked its own evaluative logic, and prompted congruence work that either reconfigured the representation or physically transformed cows and herds depending on which system was more plastic. **Figure 3.4** presents an abstract model of the representing process synthesized from these findings. The model identifies three coupled elements that operated at each moment of representing. First, distinct moments in the representing process – the inscription bottleneck and the biographical, algorithmic, and comparative renderings – each produced a different version of the cow and made different aspects of her visible. Second, each moment generated evaluative restructuring, raising the question of what was required of the physical referent to align with its digital representation and, in answering it, redefining what counted as good, desirable, and representable in the physical world. Third, each evaluative restructuring prompted congruence work aimed at resolving the gap between representational ideals and physical realities, with the direction of that work – editing the representation or editing the physical referent – determined by which of the two systems was more plastic in the moment. Together, these elements describe a recursive process in which representing continuously redefines the represented object, evaluation continuously shifts in response, and congruence work continuously reshapes either the digital or the physical side of the relationship.

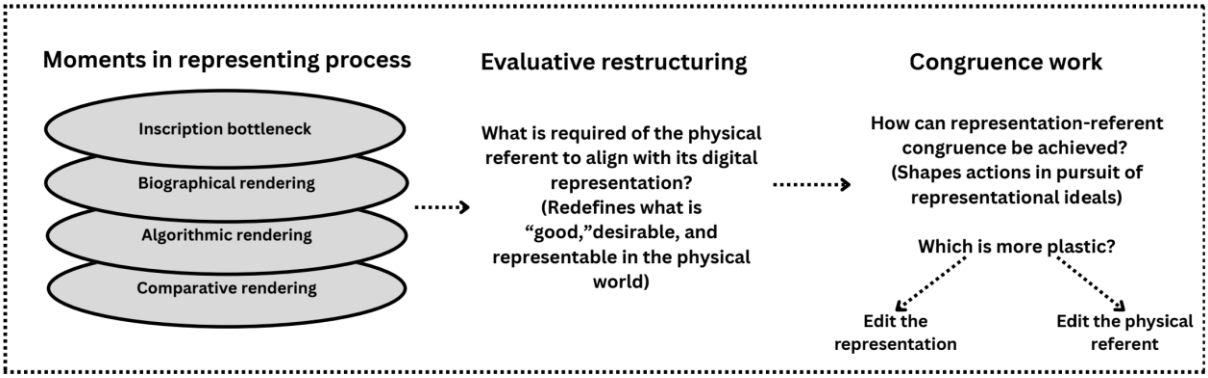


Figure 3.4. Evaluative restructuring and congruence work via evaluative moments in the representing process.

Three interrelated theoretical contributions emerge from these findings. First, I reconceptualize digital representation as an ongoing process of representing rather than a stable artifact. Second, I show how this process produces continuous evaluative restructuring that redefines what counts as a worthwhile object. Third, I demonstrate how congruence work that farmers engaged in between representations and their referents drives an ontological co-evolution in which both digital and physical entities are reciprocally transformed over time

Representing as Ongoing Process

The dominant framing in information systems and organization studies treats digital representations as relatively stable artifacts that, once produced, exert influence on cognition, decision-making, and organizational practice (Alaimo & Kallinikos, 2021; Bailey et al., 2012; Baskerville et al., 2020; Kallinikos et al., 2013; Leonardi, 2011). This framing implicitly assumes a temporal sequence wherein representations are constructed, and then consequential things happen. The question animating much of this literature is: *what do representations do?* These findings extend this line of inquiry by drawing attention to an earlier moment in the sequence – the work of *representing* itself – where much of the consequential action takes place before any stable artifact exists to be acted upon.

Observations across four moments of the representing process reveal that the construction of representations itself is where much of the organizationally meaningful action occurs – where evaluative criteria shift, where physical bodies are selected and modified, and where digital representations are curated, manipulated, and contested. The "representation" as a coherent, stable object never truly arrives. What I observed instead is the act of *representing*: an ongoing, distributed, recursive process through which both representations and their physical referents are continuously produced and revised. The cow was simultaneously rendered as a body, a data profile, a predicted trajectory, and a relative position, and none of these renderings ever stabilized into a finished depiction. The farmer's daily practice was not to consult "the representation" of a cow but to move between renderings, synthesizing, prioritizing, or adjudicating among the different versions of the cow each one offered.

This shift parallels Orlikowski's (2002) reconceptualization of knowledge as knowing – from a possessed object to a situated, enacted practice. Just as organizations do not "have" knowledge so much as they "do" knowing, representational systems do not produce representations so much as they sustain processes of representing through which both digital and physical realities are continuously reconstructed. Theoretically, this move carries three implications. First, it relocates the analytical object from the stable artifact to the practices, infrastructures, and interventions that continuously produce it: what a representational system *is* cannot be specified independently of how it is being enacted at any given moment. Second, it dissolves the assumed temporal sequence in which representations are first constructed and then exert effects – under a representing framing, construction and consequence are simultaneous, because each rendering both depicts and reshapes the referent it engages. Third,

it makes representational work continuous with organizational work more broadly, rather than a discrete technical activity that produces inputs for downstream decisions.

This processual reconceptualization has important implications for how scholars study digital representations. Where existing studies have productively examined how representations serve as boundary objects (Alaimo & Kallinikos, 2021; Carlile, 2004), coordinate work (Bechky, 2003), or shape strategic trajectories (Boland et al., 2007; Kaplan, 2011), these analyses tend to begin after the representational object has stabilized, or at least treat its stabilization as analytically prior to its effects. In contemporary data-intensive environments, these findings suggest this stabilization never occurs. The representational systems farmers used were continuous sensing, processing, predicting, and ranking infrastructures that generated new renderings with every milking, every sensor reading, and every algorithmic recommendation. Under these conditions, research designs that treat representations as stable mediating artifacts may miss the mechanisms through which digital technologies reshape organizational realities. The representing perspective calls for greater attention to the infrastructural, material, and practical conditions under which representations are continuously produced and to the recursive effects this production process has on the physical systems it engages.

Evaluative Restructuring and the Ongoing Redefinition of "Goodness"

Our second contribution concerns the evaluative consequences of representing. In the literature review, I identified a gap between the STS insight that representing is constitutive and the organizational literature's inattention to how representing reshapes the physical objects being represented. Findings address this gap by revealing how the representing process continuously generates new evaluative frameworks as a byproduct of rendering, such that the

act of representing a physical entity simultaneously multiplies and varies the criteria by which it will be judged.

At each moment in the representing process, the cow was reconstituted as a different evaluative object. At the inscription bottleneck, she became a body evaluated for device compatibility, and "goodness" meant morphological and dispositional fit. In the biographical rendering, she became a longitudinal record, and "goodness" meant an unproblematic history. In the algorithmic rendering, she became a modeled prediction, and "goodness" meant favorable projected futures. In the comparative rendering, she became a position within a distribution, and "goodness" meant invisibility (seamless and unoffensive conformity) or outperforming her peers. Each rendering brought into existence a new way of evaluating that was impossible prior to that rendering being produced.

This finding extends and complicates existing theorizations of how digital technologies reshape evaluation. Prior work has shown that quantification, commensuration, and categorization create new bases for judgment (Espeland & Stevens, 1998; Porter, 1995; Orlikowski & Scott, 2014), and that representational systems can reconfigure evaluative authority (Christin, 2020; Kellogg et al., 2020). This study contributes to this literature by demonstrating that representational systems are not mono-evaluative. The same system simultaneously sustains multiple, incommensurable evaluative frameworks because it renders the same referent in multiple ways. And these frameworks do not resolve into a hierarchy: there is no master evaluation that synthesizes morphological compatibility, genomic potential, and relative herd position into a single verdict. Instead, farmers live inside the tension between renderings, and different decision contexts activate different evaluative frameworks.

This is what makes the "invisible cow" so theoretically interesting. She is not the cow who scores highest on any single evaluative framework. She is the cow who generates no friction across any rendering – the cow for whom all four evaluative frameworks converge on the same verdict. Her invisibility is the absence of inter-rendering tension. Conversely, the most puzzling cows are those who produce divergent evaluations across renderings. Daniel's cow who produces 94 pounds of milk but will not visit the robot is difficult precisely because her biographical rendering would indicate she is an "excellent cow" while her inscription rendering suggests she is a "problem cow." The farmer's congruence work in that case is not just about aligning representation and reality, but it is also about reconciling disparate renderings that disagree about the "goodness" of that cow.

This observation points to a third contribution of this study: representing reframes congruence work from a bilateral problem (representation versus reality) to a multilateral one (multiple renderings in tension with each other and with physical reality simultaneously). The existing literature on representational fidelity treats congruence as a two-party negotiation: actors work to keep the representation aligned with the referent, or *vice versa*, and the analytical interest lies in which side gets edited (Burton-Jones & Grange, 2013; Bailey et al., 2012; Østerlie & Monteiro, 2020). Findings show that this framing understates the problem. When multiple renderings of the same entity coexist, congruence work involves at least three distinct relationships at once: each rendering must be aligned with the physical referent, but the renderings must also be reconciled with one another, and the priority a farmer assigns to one rendering over another reshapes the criteria by which the others are subsequently judged. Daniel did not simply ask whether his cow's data was accurate; he had to decide which version of her counted, and what to do when those versions disagreed.

This multilateral structure introduces forms of work that the bilateral framing renders invisible. Farmers spent considerable effort adjudicating between renderings (deciding which framing to act on in a given moment), sequencing them (treating one rendering as the standing default and another as a check), and insulating them from one another (e.g., manipulating thresholds in one rendering so its outputs would not contaminate the evaluative criteria in another). These forms of work are not failures of fidelity but routine features of operating inside a multi-rendering representational system, and they have no analogue in accounts that treat representations as singular artifacts in correspondence with referents. The contribution, then, is not only to identify that congruence is multilateral, but to specify the forms of organizational labor that multilateral congruence requires.

The theoretical significance of this observation lies in what it reveals about the relationship between ways of knowing and the material world. The literature on representational decoupling has warned of "liquefaction" (Monteiro & Parmiggiani, 2019) or "ontological reversal" (Baskerville et al., 2020) in which representations precede and become independent from the physical object they represent. Findings reveal a different dynamic. In organizational contexts that continuously monitor active, dynamic physical systems, decoupling is unlikely, if not impossible. Instead, the epistemic and ontological are inextricably entwined. The process of representing did not merely change how farmers know cows; it changed which cows existed, what their bodies looked like, and how herds were composed. When morphological incompatibility with robotic milkers led to culling, when genomic scores directed selective breeding, and when distributional rankings determined which genetic lines persisted, the epistemic apparatus of representing reached directly into the material constitution of its referents. New ways of knowing produce new physical realities. This dynamic resonates

with Mol's (2002) insight that practices of measurement enact their objects, Leonardi's (2012) demonstration that virtual representations reshape material design, and Goodwin's (1995) observation that instruments constitute the objects they render knowable. This study extends these insights by showing how this constitutive relationship operates through multiple concurrent renderings whose evaluative pressures accumulate and compound, producing gradual, incremental, and often imperceptible co-evolution of digital and physical realities over time.

Congruence Work and the Materialization of the Epistemic

Our third contribution theorizes congruence work as the mechanism through which the epistemic consequences of digital representation materialize into physical action and change. The literature on digital technologies has documented how representational systems shape what people know and how they know it (Orlikowski & Scott, 2014; Zuboff, 1988). What has received less attention is how these shifts in knowing translate into concrete, often irreversible, physical changes and what governs the direction of that transformation. Findings demonstrate what is often piecemeal in the literature: the recursive dynamic in which representing shapes evaluation, evaluation shapes intervention, and intervention reshapes the object that is then re-represented. Findings show that congruence work is a constitutive practice through which the epistemic and the material become entangled.

Throughout the representing process, farmers acted as mediators between ontological domains. They interpreted data on screens while physically examining cows, reconciled predictions with embodied experience, and balanced comparative rankings against contextual knowledge. Congruence often required judgment, discretion, and sometimes defiance of what either the representational system or their own senses were telling them. Yet resistance operated

within the constraints of the representational technologies. Whether overriding predictions or acting on them, farmers remained oriented toward maintaining correspondence between what representations depicted and what herds physically manifested.

Critically, the direction of congruence work – whether farmers edited the representation or the referent – was governed by the relative ontological plasticity of the two systems. When representational systems were malleable (thresholds adjustable, data records editable), farmers modified the digital. When representational systems were immutable or opaque (hardware constraints, proprietary algorithms, aggregate metrics), farmers modified the physical. This plasticity-dependent directionality means that the design of representational systems – specifically, which elements are made adjustable and which are fixed – has direct consequences for whether congruence work reshapes the digital world or the physical one.

Although these findings emerged from dairy farming, the plasticity-dependent directionality of congruence work is likely a general feature of any organizational setting in which digital representations are continuously reconciled with physical, embodied, or institutional referents. In hospitals, clinicians routinely encounter discrepancies between the patient as rendered in the electronic health record and the patient in front of them: when the EHR's diagnostic categories or order sets are rigid (locked clinical pathways, hard stops in the medication module), clinicians adjust their documentation, ordering practices, or even patient classification to fit the system (Bossen & Markussen, 2010); when the record is malleable, they annotate, override, or work around it to preserve the clinical reality. In platform-mediated labor markets such as ridesharing, food delivery, or freelance gig work, drivers and workers cannot edit the algorithmic representations that score and rank them, so they modify their physical behavior (e.g., choosing routes, timing breaks, gaming acceptance rates) to produce the digital

trace the system rewards (Cameron, 2022). In warehouses and distribution centers, when scanner-based tracking systems impose immutable pick rates, workers reorganize their bodies and movements to meet the metric; when the system tolerates manual overrides, supervisors edit the digital record to absorb physical variability (Delfanti, 2021). In policing, predictive systems and dashboards generate rankings that officers cannot easily alter, prompting interventions in the physical world (e.g., patrol patterns, stop-and-frisk practices, arrests) that bring the territory into alignment with the map (Brayne, 2021). Across all these domains, the same logic appears: the side that yields is the side that can yield, and the design choices that determine plasticity (open data fields versus locked schemas, adjustable thresholds versus proprietary models, editable records versus immutable logs) directly shape whether the consequences of representational systems accrue in the digital or the physical world.

This observation carries significant theoretical and practical weight. It means that the architecture of representational systems channels the direction of physical action. When a robotic milker's arm could not be reconfigured, the cow had to be molded to adapt. When an algorithmic threshold was gatekept by a vendor, the animal population was adjusted to align with predictions. When herd-level averages could not be directly edited, individual cows were modified on a genetic level to shift the aggregate. In each case, the rigidity of the representational infrastructure redirected congruence work toward the physical realm. This extends Barley's (1996) concept of the "empirical interface" at which technicians steward both physical and representational systems, by showing how the relative malleability of each side of that interface determines the material trajectory of organizational action.

What makes this process particularly consequential is its recursive and compounding nature. Impacts of congruence work at one rendering can "leak" into other renderings via

altering the representational material on which they are built. Downstream, the cumulative effects can be invisible even to the actors driving them. When farmers culled cows whose bodies did not fit robotic milkers, they altered the herd composition that generated the data visible in the comparative rendering. When they treated cows based on algorithmic predictions that included false positives, they introduced treatment records that became part of the cow's digital history and influenced future evaluations. When they bred selectively based on distributional rankings, they shifted the herd average against which all future cows were measured. Each local act of congruence work – pragmatic, situated, and unremarkable in isolation – fed into a cascade of downstream consequences that altered the conditions under which subsequent rounds of representing occurred. Over time, the accumulation of these decisions guided representation-referent co-evolution through the gradual, concurrent transformation of digital and physical realities that no single decision intended but that the structure of the representing process systematically generated.

The long-term consequences of this process may surprise even the actors who drive them. Farmers who adjusted pregnancy detection thresholds did not anticipate the cascade of false-positive abortions that followed. Farmers who selected for robotic compatibility did not set out to narrow the morphological diversity of their herds, yet cumulative compatibility-based culling produced exactly this outcome. The rising genetic baselines that farmers celebrated as progress simultaneously raised the bar for "good enough," creating intensifying selection pressures on animals that would have been adequate in a prior context. These unintended consequences are structural products of the representing process – emergent properties of a system in which each rendering generates its own evaluative demands, congruence work resolves immediate misalignments without full awareness of accumulating impacts, and the

physical referents that survive this process become the baseline against which subsequent renderings are produced.

Conclusion

At its core, the process is a story of domestication and artificial selection – the deliberate modification of organisms to serve human needs – but with a twist. Selection pressures are typically thought of as stemming from human agency (e.g., farmers breeding for desirable traits, breeders selecting for specific phenotypes). What these findings reveal is a less obvious force guiding the direction of selection: evaluative restructuring embedded in the representing process and the technological agency imbued throughout. By redefining what is considered good and desirable through successive renderings, representational technologies shape what organizations select for or against. Throughout the process of representing a cow, farmers were forced to make choices – how to interpret representations, where to direct congruence work – that reconfigured the literal bodies of cows according to evaluative criteria introduced by the representational system. The notion that technology shapes how we cultivate the natural world is not new. But as representational technologies become more pervasive, more algorithmically prescriptive, and more capable of rendering complex phenomena across multiple platforms and formats, we are witnessing a subtle shift in the locus of ontological authority that could have significant implications for how physical realities (human, animal, and environmental) are structured.

These insights extend well beyond dairy farming. As representational technologies proliferate across healthcare, urban planning, environmental management, and other domains where dynamic digital systems engage dynamic physical systems, the mechanisms I identify are likely to operate wherever representing is continuous and the referents are physically

malleable. The plasticity-dependent directionality of congruence work suggests that the rigidity or flexibility built into representational infrastructures shapes whether misalignments are resolved by adjusting the model or adjusting the physical world. In domains where the referents are ecosystems, bodies, or communities, the stakes of this directionality are profound. Understanding how the process of representing reshapes the physical world is essential for designing representational systems that account for, rather than obscure, their constitutive role in shaping the realities they purport to describe. It is through examining this entanglement that we can clearly observe the way the hammer shapes the hand.

CHAPTER 4: Performing as a Recursive Process of Orienting – How Pre-Existing Performance Threats Shape Technology Use in Practice

Experts are increasingly faced with new technologies that disrupt how they work (Barley, 2020; Faraj et al., 2018; von Krogh, 2018), which requires them to make strategic choices about how to engage with technology while maintaining their expert standing (Bailey & Leonardi, 2015; Van Den Broek et al., 2021). Scholars have examined the relationship between technology and expertise by studying how technologies can replace or supplement the knowledge experts use in their work (Acemoglu & Restrepo, 2019; Dell'Acqua et al., 2023; Frey & Osborne, 2017; Raisch & Krakowski, 2021), how experts make sense of technological disruption (Nelson & Irwin, 2014; Van Doorn et al., 2023), and how an expert's assessment of a new technology's capabilities shapes their subsequent use decisions (Allen & Choudhury, 2022; Lebovitz et al., 2022; Orlikowski & Gash, 1994). How experts interpret and respond to technological change can have serious implications for the success of technology adoption (Edmondson et al., 2001), as well as for the development and maintenance of expertise more broadly (Anthony, 2021; Beane, 2019; Leonardi, 2007). The resulting image, proffered by numerous empirical studies, is that advanced technologies threaten the roles played by experts (e.g., Barley, 1986; Beane & Anthony, 2023; Bechky, 2020; Christin, 2017; Orlikowski, 2000; Valentine & Hinds, 2022). Although technologies can surely also provide opportunities to enhance expert performance for some (Jia et al., 2024; Prasad, 1993; Treem, 2013; Waardenburg et al., 2022), threat remains the dominant watchword in studies of expertise and technology use.

Understanding how new technologies threaten the standing of experts is particularly challenging due to two problematic assumptions in the existing literature. First, most studies treat expertise as existing in relative equilibrium before technological disruption, viewing

experts as having established recognition before technology challenges their position (Barley, 1990; Barrett et al., 2012; Beane & Anthony, 2023). From this perspective, experts respond to technology – induced threats through new skill development (Beane, 2019), changing their practices (Leonardi, 2007; Sergeeva et al., 2020), and adopting new thinking (Fayard et al., 2016; A. Nelson et al., 2023; Pachidi et al., 2021) to preserve status and identity (Lifshitz-Assaf, 2018; Vaast & Pinsonneault, 2021). However, this view underemphasizes how expertise is continuously performed, negotiated, and contested through ongoing interactions (Chan & Hedden, 2021; DiBenigno, 2020; Huising, 2015; Treem & Leonardi, 2017), making expert standing already tenuous before technologies arrive into contexts where experts manage existing challenges (Bourgoin & Harvey, 2018; Gieryn, 1983; Kang & Kim, 2022; Treem, 2016). Second, literature assumes expertise is threatened as a direct consequence of technology's arrival, with responses emerging only after engagement (Barley, 1990; Lebovitz et al., 2022; Mazmanian, 2013). This temporal bias de-emphasizes how current thinking results from ongoing expertise maintenance predating technology's arrival and how members interpret technologies before implementation (Leavell, 2023). This "left-censored" view (Leonardi & Barley, 2010, p. 38) misses how expert responses extend pre-existing expertise performance work addressing threats that existed before implementation. In short, rather than thinking of technological change as an exogenous force that threatens to disrupt the calm waters of expertise, we might think of expertise as a continuously flowing and evolving river into which technology enters and becomes incorporated into the flow.

To address these limitations, I propose that expertise should be understood as a dynamic, relational, and context-dependent process (Abbott, 1988; Ericsson et al., 2018; Pakarinen & Huising, 2023) in which technologies serve as 'props' that might be used to

stabilize or disrupt ongoing expertise performances (Pinch, 2010). This perspective shifts attention from how technology directly threatens the work of experts to how it becomes incorporated into the ongoing work of expertise maintenance in which experts are continuously engaged. It also emphasizes that understanding technology's impact requires examining the pre-existing expertise dynamics that shape how experts initially make sense of and engage with new technologies (Orlikowski & Gash, 1994). By integrating an "expertise-as-performance" lens (Eyal, 2013; Goffman, 1956; Sandberg et al., 2017; Treem, 2016) with insights from the literature on threat discernment (Jackson & Dutton, 1988), I develop a more dynamic, situated understanding of the expert-technology relationship more befitting our rapidly evolving technological world of work.

To demonstrate the utility of this approach, I draw on a field study of dairy farms implementing AI-powered robotic milking systems that reveals how individuals perceive the same technology differently depending on the precariousness of their existing expertise performances. I demonstrate how a *lens of threat*, which is shaped by the intensity and specificity of threats to ongoing expertise performance experienced before new technologies arrive, influences how experts interpret new technologies and shapes their engagement with them. I find that experts who experience ongoing threats to their standing use new technologies strategically to address existing vulnerabilities in their expertise performance, hoping to preserve or reclaim their expert status. I discuss the implications of our findings for studies of technological change and expert work, as well as for expertise more broadly.

Perspectives on Technological Change in Expert Work

Studies of technological change and expertise have repeatedly documented that novel technologies can disrupt expert work. As the technologies of interest have evolved from basic

computerized systems for record keeping and imaging (Barley, 1986; Orlikowski, 2000; Robey, 1979; Zuboff, 1988) to more advanced artificial intelligence systems designed to automate or augment decisions (Anthony, 2021; Endacott & Leonardi, 2022; Jarrahi, 2019; Lebovitz et al., 2022) and robotic systems that reshape who does work and what kind of work is done (Barrett et al., 2012; Beane, 2019; Sergeeva et al., 2020), technological change is consistently depicted to be a force that threatens the stability and standing of experts. **Table 4.1** summarizes numerous influential studies that have documented and explicitly discussed how new technologies threaten the work of experts. The trenchant analyses in these studies and others reveal that technological threats to expertise operate through four primary mechanisms: (1) redefining the knowledge and cognitive foundations of expert work, (2) disrupting established practices and routines, (3) undermining expert identity and self-concept, and (4) reconfiguring relational and structural boundaries.

Table 4.1. Summary of literature on technological threats and expert response.

Paper	**Technology as a threat**	**Expert responses**
Anthony, 2021	"Surprisingly, those who perceived technologies as potentially **threatening** did not avoid the technologies or reframe their own expertise but rather engaged more deeply with the assumptions in their tools through task allocations and oversight than those who perceived them as marginally affecting work." p.30	In one group, senior and junior members split tasks in ways that isolated each group from key learnings. In the other group, active oversight and reasoning enforcement by senior experts prevented black boxing.
Barley, 1986	"Clandestine teaching **threatened** the institutionalized roles of radiologists and technologists. Under radiology's traditional system, radiologists taught technologists, but the reverse was uncommon and nearly taboo. Only the radiologist's front of self-assurance and the technologist's deference, both of which were encoded in the semantics and pragmatics of the exchange, kept clandestine teaching from becoming open instances of	Different radiology teams developed distinct "scripts" that shaped how work was distributed.

	role reversal." p.92	
Beane, 2019	"For the resident, this meant a dramatic reduction in opportunities to ''do'' parts of the procedure without significant supervision and a near elimination of opportunities to ''teach'' other more junior residents. These issues alone presented a significant **threat** to residents' acquiring robotic surgical skill through legitimate peripheral participation, but these dynamics were compounded by the way that APs supervised and taught residents during robotic surgical work." p.13	Surgery residents had to resort to deviant learning methods to obtain the necessary skills to perform their work.
Beane and Anthony, 2023	"These individuals were supposed to know what to do and to look like it: engaging in the visible struggle traditionally associated with learning could **threaten** all of this." p.9	Senior and junior members adopted "inverted apprenticeships" that reconfigured relationships where seniors learned from juniors while maintaining their authority.
Bechky, 2020	"The institutional pathway by which toxicologists encountered an evaluative comparison was a defense attorney's immediate challenge to their presentation of specific case results in court in which he would draw a parallel to error measurement techniques used in DNA profiling. It was a direct **threat** to their ability to successfully testify rather than a broader legitimacy challenge such as that to the firearms examiners' technique." p.627	Varied by occupation: rejection, adaptation, or selective resistance based on value alignment.
Christin, 2017	"Moira explains that her main role was to make the journalists use web analytics: 'My mandate was to get more traffic… they didn't have a clear strategy about how to do it. When I arrived people talked very differently about traffic, they saw it as **threatening**, debasing. That's not the case anymore.'" p.8	Experts engaged in "buffering strategies" including passive resistance, strategic manipulation, and direct opposition to preserve control over decisions.
Geampana & Perrotta, 2023	"To a certain extent, the existence of TL and its functions **threaten** to disrupt the current order: the technology exposes patients to what was largely "inside knowledge" before. As such, it becomes evident that underneath the surface of TL's algorithmic promises, lie uncertainties regarding the best use of this technology and information generated through it." p.227	Embryologists responded to the algorithm by strategically integrating its output into their decision-making while reaffirming their expert judgment, selectively deferring to, overriding, or critiquing the model to

		preserve their epistemic authority and moral responsibility.
Groleau, Demers, Lalancette, & Barros, 2012	"At the end of the study, architects expected that the institutionalization of practices favoring the production of realistic visuals would **threaten** the recognition of their exclusive expertise by clients and thus challenge rules embedded in the regulatory context." p.15	Experts reluctantly adapted to preserve commercial viability while symbolically defending professional values.
Jia, Luo, Fang, & Liao, 2024	"While highly skilled agents described mixed feelings of appreciating the help they received from AI and concern about being replaced by AI in the future, they explicitly said that the company had adopted the right strategy. For example, 'In the short term, [AI assistance] is a good thing, but in the long run, there will be **threats**.'" p.25	High-skilled experts embraced enhancement while lower-skilled workers sought support and resisted expansion.
Lifshitz – Assaf, 2018	"This analysis revealed a clear hero figure: the ''innovator,'' whose professional identity was that of a problem solver. Some R&D professionals saw the open innovation model as a **threat** to this identity, while others discussed the need to change their identity in response to the potential of the open innovation model from being problem solvers to being solution seekers." p.753	Divergent identity-based responses: boundary dismantling through identity refocusing vs. boundary protection and feigned adoption.
Nelson & Irwin, 2014	"This period witnessed the introduction of Yahoo, which would go on to become the decade's most popular search engine, and which, critically, employed an approach to searching that mimicked the subject categories at the heart of librarians' traditional approach to searches. In this similarity, however, librarians perceived a **threat**." p.908	Sequential identity transformation to redefine expertise – from gatekeeper to educator to connector.
Orlikowski, 2000	"The third reason consultants enacted a limited – use technology-in-practice with Notes arose from their fear that use of its collaborative properties would **threaten** their status within Alpha." p.417	Experts responded in divergent ways – ranging from minimal or skeptical use, to leveraging it for individual efficiency, to creatively improvising and integrating it into collaborative knowledge-sharing practices.
Scarbrough, Chen, & Patriotta, 2024	"For this grouping, concerns about the potential loss of control of their tasks and the reduction of their professional expertise meant that AI was perceived as a **threatening** incursion into their professional	Role-differentiated sensemaking: front-line resistance, managerial pragmatism, and

	jurisdiction – in effect an 'invading actor' in their work environment. In contrast, the Manager grouping viewed the implications of AI in terms of opportunities for increased system efficiency, and the need to integrate AI into the organization as an instrument of such efficiency. AI was seen as decoupled from, rather than **threatening**, their jurisdiction." p.24	leadership reframing as opportunity.
Vaast & Pinsonneault, 2021	"We offer original insights in this regard by explaining how the generativity of digital technologies leads to an occupational identity that is closely associated with IT identity but also with the occupational identity that is continuously **threatened** by obsolescence through digital technologies." p.1103	Ongoing identity negotiation through claims of interdisciplinary value, strategic insight, and balancing threat of obsolescence with finding distinctiveness.
Valentine & Hinds, 2022	"We analyze the data scientists' and buyers' interactions that produced the theory testing capabilities and suggest that testing knowledge workers' theories using algorithms is a new, co-produced expertise that both builds on and **threatens** their expertise. p.144	Collaborative co-production of new "theory-testing" expertise through human-AI integration, adaptation, and skill development.
Yam, Tang, Jackson, Su, & Gray, 2023	"our multimethod and multicultural findings specifically demonstrate that the adoption of robots can perhaps be equally influential in inducing job insecurity among employees working in both low skilled and intellectually demand jobs. Thus, our work sheds further lights on the consistently **threatening** nature of robots." p.10	Defensive psychological responses are marked by job insecurity, burnout, and workplace incivility.

First, technology redefines cognitive foundations by digitizing tacit knowledge and transforming situated understanding into standardized algorithmic inputs, stripping away contextual nuances experts rely upon for judgment (Alaimo & Kallinikos, 2022; Brayne, 2017; Zuboff, 1988). Lebovitz and colleagues (2021) show how diagnostic AI models codify expert "know-what" without incorporating contextualized "know-how" in hospitals. Second, technology disrupts practices by requiring new technical skills (Beane, 2019), reconfiguring routines (Edmondson et al., 2001), and increasing visibility of previously inscrutable actions (Karunakaran, 2022; Leonardi & Treem, 2020). Beane and Anthony (2023) demonstrate how

senior bankers and surgeons adopted "inverted apprenticeships" with juniors to update technological expertise while maintaining positional status. Third, technology undermines identity as technologies prompt experts and evaluators to question the value and primacy of their work (Bechky, 2020; Lifshitz-Assaf, 2018). Nelson and Irwin (2014) show librarians successfully incorporated internet search only after redefining their occupational identity to align with new technological capabilities. Finally, technology reconfigures relational structures by redistributing knowledge across roles and creating hybrid practices that blur jurisdictions (Van Den Broek et al., 2021; Wang et al., 2024). Barley's (1986, 1990) work on CT scanner implementation demonstrates how technology-induced task changes reverberate through social systems, altering roles and reconfiguring expert hierarchies. **Table 4.2** explores these threats and their mechanisms in greater depth.

Table 4.2. Groups of mechanisms through which technology disrupts expertise.

Categories of disruption	Representative mechanisms	Examples in literature
Disrupts knowledge and cognition (cognitive)	Datafication and decontextualization of information	Alaimo & Kallinikos, 2022; Bailey et al., 2012; Brayne, 2017; Zuboff, 1988
	Automates tasks or codifies knowledge, deskilling, down skilling	Acemoglu & Restrepo, 2019; Braverman, 1998; Brynjolfsson & McAfee, 2014; Dell'Acqua et al., 2023; Frey & Osborne, 2017
	Disrupts cognition, decision-making, and judgment	Lebovitz et al., 2022; Lindebaum et al., 2020; Valentine & Hinds, 2022
	Changes the knowledge bases and tools experts must engage with	Beane, 2019; Grønsund & Aanestad, 2020
Disrupts practices and routines (practical)	Requires new skills, practices, and proficiency with technology	Abdel – Karim et al., 2023; Anthony, 2021; Cristea & Leonardi, 2019; Leonardi, 2007
	Reconfigures roles and routines	Barrett et al., 2012; Beane & Anthony, 2023; Christin, 2017; Edmondson et al., 2001; Lifshitz – Assaf, 2018

	Increases visibility and accountability	Karunakaran et al., 2022; Leonardi & Treem, 2020
Disrupts identity and self-concept (symbolic)	Affects self-concept, identity, and self-presentation	Davis & Hufnagel, 2007; Leonardi & Treem, 2012; Nelson & Irwin, 2014; Vaast & Pinsonneault, 2021
	Changing evaluation metrics, new benchmarks, symbolic redefinition of expertise	Bechky, 2020; Mazmanian & Beckman, 2018
	Erosion of authority, jurisdiction, autonomy, dignity, and discretion	Faulconbridge et al., 2023; Galperin, 2020; Lifshitz – Assaf, 2018; Scarbrough et al., 2024; Zuboff, 2015
	Evokes doubt, anxiety, and technostress	Ayyagari et al., 2011; Fügener et al., 2021; Yam et al., 2023
Disrupts relations and structure (structural)	Redistributes knowledge, changes knowledge structures, while benefits accrue unevenly across roles/occupations	Barley, 1986; Chen & Chan, 2024; Leonardi, 2007; Wang et al., 2024; Wu & Kane, 2021
	Changes public perception of expert groups	Bechky, 2020; Kronblad, 2020
	Alters boundaries and creates new gatekeeping, brokerage, and hybrid/ambiguous roles	Van Den Broek et al., 2021; Waardenburg et al., 2022

Of course, across contexts, there are a wide range of expert responses to technological change. Some experts embrace new technologies as opportunities for professional enhancement (Prasad, 1993), others resist them as existential threats even to their own detriment (Allen & Choudhury, 2022; Kang & Kim, 2022), and still others develop hybrid approaches that selectively incorporate technological capabilities that enhance expertise while rejecting those perceived as threats, sometimes producing unanticipated and unintended outcomes for both experts and organizations (Pachidi et al., 2021). Despite the vast literature, current approaches may be missing key dynamics that shape how experts interpret and engage with new technologies (Bailey & Barley, 2020). While studies vary in context and focus, two consistent assumptions limit our understanding of the expertise-technology relationship. First,

the literature assumes that technology itself serves as the primary source of threat to expertise, and second, that experts respond reactively to technological threats as their interpretations emerge from interactions with these tools in the practice of use. Each of these assumptions, although reasonable in nature, impose limitations on how scholars approach and theorize phenomena around technological change and expertise.

Technology as a Threat to Expert Stability

One important assumption that permeates the technology and expertise literature is that expertise exists in an established state of equilibrium when new technologies arrive. Studies consistently begin their analysis by characterizing experts as securely positioned in their roles, with recognized knowledge, established authority, and uncontested legitimacy before technological disruption occurs (Barley, 1986; Beane & Anthony, 2023; Bechky, 2020). For instance, Edmondson and colleagues (2001) demonstrated how a new surgical tool altered the established role boundaries of operating room teams in ways that disrupted the routines that reinforced institutionalized hierarchies. Similarly, Lebovitz and colleagues (2022) described radiologists as practicing uncontested diagnoses before AI systems injected uncertainty into their decision-making processes.

This focus on expertise as stable inadvertently misses the complex pre-histories of expertise performances that precede a technology's arrival, as well as the nuanced social dynamics that shape how experts come to interpret these technologies. Bounded by empirical scope, scholars observe expert practices and consequences that emerge from interactions with the tool but regularly miss out on potential insights from informants' accrued experiences that might help explain how they interpret new technology (for a more robust discussion of this problem see Bailey & Barley, 2020; Leonardi, 2009). **Figure 4.1** depicts how this narrow

temporal aperture can lead to a logical but potentially spurious diagnosis of expert stability before the implementation of a new technology. Expertise is volatile. Experts are constantly performing their expertise for others while audiences have made varying attributions of expert standing based on these performances. But most organizational scholars who adopt a pre-/post-implementation study design do so to observe changes to expert work after a new technology is introduced (Barley, 1986; Leonardi, 2013). This approach may just be capturing a relatively small period of expert performances that appear stable, when it is really just a moment in an on-going process. However, the conditions under which technology implementation is viewed positions it to be seen as a discrete, exogenous shock that enters otherwise presumed stable conditions.

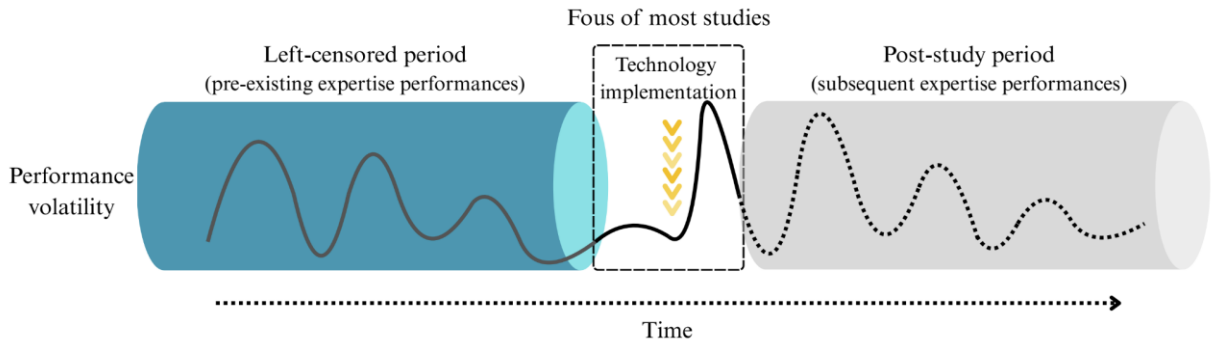


Figure 4.1. Studies of expertise in relation to the broader temporal context of expertise performance.

This omission of seeing implementation of a new technology as but a moment in a longer performative arc is what Leonardi and Barley (2010, p. 38) call a "left-censored" view of technological change that "makes it impossible to determine whether patterns of use are shaped in important ways by dynamics of power, control, status, and conflict that set the context of use." Even if expertise were stable at the time of implementation, neglecting experts' past experiences forfeits a vast wealth of data about formative encounters, drawn out negotiations, and often consequential interactions that form the basis of how experts perceive,

interpret, and respond to subsequent events (Bar-Haim et al., 2007; Floyd et al., 2000; Staw et al., 1981). This bias leads scholars to overlook the dynamic social processes through which expertise is continuously claimed, negotiated, and maintained (Bourgoin et al., 2020; Chan & Hedden, 2021; Suddaby & Greenwood, 2005; Treem, 2016). By extending our analytical scope to include more nuanced portrayals of pre-implementation contexts, scholars can enrich existing theories by recognizing how past experiences and ongoing expertise performance shape how experts respond to new technologies.

A performance approach, grounded in Goffman's (1956) dramaturgical perspective and Eyal's (2013) sociology of expertise, overcomes this stability assumption by recognizing expertise as a practical, relational, and situated exercise rather than a fixed state. Goffman argues that all social identities must be continuously performed for audiences through careful management of face – an "image of self, delineated in terms of approved social attributes" (Goffman, 1967, p. 5). Building on this insight, Pinch (2010, p. 414) postulates that technologies play a part in "staging" of roles and "mediation" of interactions in these Goffmanian performances, participating in the ongoing work of constructing and maintaining expert identities rather than simply disrupting them from offstage. Moreover, sociologists like Eyal (2013) and Abbott (1988) argue that expert occupations require participation from other groups and continuous negotiations of boundaries and jurisdictions in order to constitute expertise. Contemporary expertise scholars have further developed this performance perspective, emphasizing how expertise must be continuously demonstrated and validated through social interactions (Collins & Evans, 2007), how experts use face to influence patterns of interaction while maintaining coherent identity (Patriotta & Spedale, 2009), how recognition depends on ongoing congruence between self-perceptions and others' perceptions of

competence (Treem & Leonardi, 2017), and how expertise is embedded in relations (Pakarinen & Huising, 2023).

Viewing expertise as an ongoing performance rather than a stable state reveals that technologies enter dynamic contexts where experts are already managing various challenges to their standing, rather than disrupting otherwise stable arrangements. Studies demonstrate the inherent "fragility of experts" (Kang & Kim, 2022, p. 577): management consultants routinely face credibility threats when appearing knowledgeable while learning unfamiliar client contexts (Bourgoin & Harvey, 2018), career advisors confront legitimacy threats when client values contradict professional guidance (Chan & Hedden, 2021), and experts across fields experience jurisdictional threats from competing occupations claiming authority over their domains (Lifshitz-Assaf, 2018; Truelove & Kellogg, 2016). These ongoing challenges require continuous performance adjustments, highlighting how expert work demands persistent negotiation and adaptation even without technological change. Technology is not necessarily the primary threat; rather, it enters performances where experts are constantly defending their standing through technical, symbolic, and relational work (Carlile, 2004; DiBenigno, 2020; Huising, 2015; Vaast & Pinsonneault, 2021). From this perspective, technology's role depends on existing performance dynamics; it might bolster an already destabilized performance or exacerbate existing vulnerabilities. This performance lens helps explain why experts differ in interpreting new technologies and what engagement strategies they choose based on the specific challenges they were already navigating. Rather than assuming technology disrupts stable expertise, we should recognize how experts incorporate technology into their ongoing efforts to maintain standing, with their responses shaped more by pre-existing performance pressures than by the technology itself.

Technology as the Starting Point for Interpretation

Another assumption in the literature involves the temporal dynamics of technology interpretation. Studies typically assume that experts initially engage with new technologies in relatively neutral fashion, perhaps with some pre-emptive influence from leadership or prior experiences with similar technologies that shape their interpretations (Orlikowski & Gash, 1994; Treem et al., 2015). However, when formulating appraisals of threat or opportunity, the majority of experts' interpretations are depicted to emerge gradually through use, interaction, and experience with technological features post-implementation (Barley, 1986; Mazmanian, 2013; Orlikowski, 1996; Scarbrough et al., 2024). This gradual sense-making narrative portrays experts as developing their assessments of technological threats through extended engagement, learning about limitations and capabilities over time, and adjusting their responses accordingly (Weick, 1990). For example, Pachidi and colleagues (2021) described sales professionals as initially questioning algorithmic accuracy before eventually shifting to symbolic compliance, while Leonardi (2007) showed how emergent experiences with a new information technology service management tool led IT experts to develop novel advice seeking strategies that took advantage of the technology's capabilities. This temporal framing suggests that technology interpretation is primarily reactive, emerging from direct experience with technological capabilities and constraints.

Research on threat discernment, however, suggests that experts rapidly categorize new technologies based on existing vulnerabilities rather than through extensive interaction. Jackson and Dutton demonstrate that individuals quickly classify ambiguous stimuli as threats or opportunities based on perceived relevance to existing concerns, with managers exhibiting a consistent "threat bias" whereby they are more sensitive to threat characteristics than

opportunity signals (1988, p. 385). New technologies fit precisely the criteria for such rapid categorization: they are novel, ambiguous, and potentially consequential for expert standing. Social psychology and cognitive science reveal numerous mechanisms underlying these rapid threat appraisals. Threat perceptions trigger motivated reasoning processes that filter information according to existing concerns and goals – what we want or need shapes what we see and believe (Jost et al., 2003; Steele, 1997). The threat-rigidity response narrows attentional focus and increases reliance on dominant response patterns (Staw et al., 1981), while appraisal theory explains how individuals evaluate potential stressors in terms of personal significance and available coping resources (Lazarus & Folkman, 1984). Research shows threatened individuals react with varying sensitivities, responding with goal-directed action or avoidance to overcome threats (Gray, 1982). As Jonas and colleagues explain, "specific motivational needs for self-preservation, certainty, control, or self-esteem may guide specific reactions to particular threats under specific circumstances" (2014, p. 261). This suggests experts may evaluate new technologies through the lens of their existing performance challenges, forming threat appraisals based on pre-existing vulnerabilities rather than through an objective evaluation of technology's inherent characteristics.

If expertise is a performance that is already occurring when a new technology arrives on the scene, and if the performance is orchestrated to defend expert standing in the face of existing threats, it would behoove scholars of technology and expertise to understand how these existing performances shape the interpretation of a new technology and its incorporation into their ongoing performances. Put more formally, the central question that guides this study is: *How do experts' ongoing performances of expertise and prior experiences of threat shape*

their interpretations of new technologies, and what are the implications for how they subsequently engage with these tools over time?

Methods

Research Setting and Site Selection

To answer this question, I draw our data from a large-scale, comparative field study of twelve dairy farms in California and New York collected between August 2022 and August 2025. The project began as an inductive investigation of an understudied empirical phenomenon with a broad research question: How do new technologies and data change the practice of farming? Our team deployed a relational ethnographic methodology to observe complex interactions between entities around technology use "in the wild" to develop a grounded understanding of the intersection between technology and work in practice (Anthony et al., 2023; Bailey et al., 2022). This approach allowed patterns and theoretical insights to emerge organically from our observations.

Out of the twelve farms studied, I identified four farms in California that implemented robotic milking systems from two competing manufacturers (with one farm adopting the first and three farms adopting the other) to monitor and enhance cow care and performance in an increasingly challenging agricultural labor market. I selected these sites because they: (1) implemented virtually identical robotic milking systems within similar time periods; (2) operated under similar market, regulatory, economic, cultural, and geographic conditions; and (3) represented comparable organizational structures, activities, and operations. This selection strategy controlled for potential confounding factors while allowing me to observe variations in how farmers adapted to the onslaught of new technologies and the data they availed.

The robotic milking systems on each farm relied on advanced sensor-based devices, including automated milking robots, milk quality analyzers, activity tracking collars, and RFID tags. These were paired with AI-enabled software programs that used computer vision and pattern recognition, trained by vendors on data related to cow behavior, udder position, somatic cell count, and milk yield. Together, the systems could predict the likelihood of udder infections, systemic illness, and ovulation, while also collecting data on milk production, feeding and rumination behaviors, and health events for each cow. This system produced vast quantities of data and algorithmic predictions that altered how farmers made decisions. Traditionally, the dairy industry had relied heavily on the embodied expertise of experienced farmers and specialists to make such predictions. For example, two predictions core to dairy farming – mastitis (a potentially deadly udder infection) and heat (the stage of estrus in which a cow is fertile and ovulating) – required specialized knowledge and use of visual and manual examination of cows to accurately predict, while detecting illness often required more tacit, intuitive "hunches" about cow comfort. The new robotic milking systems provided unprecedented data that transformed expert work, most notably, in mastitis detection, illness prediction, and heat detection, where traditional tacit and embodied knowledge encountered formalized, quantified metrics about cow health, behavior, and milk production.

The introduction of robotic milking systems represents precisely the ambiguous stimulus that Jackson and Dutton (1988) identify as triggering threat categorization. First, they carry negative implications for traditional expertise – potentially displacing farmers' sensory, embodied knowledge and intuitive decision making that form the foundation of expertise. Second, they create low control and competence feelings as farmers must learn to incorporate algorithmic predictions they cannot fully understand, while proprietary software limits their

ability to modify recommendations, creating the control anxieties Jackson and Dutton associate with threat perception. Third, systems present ambiguous information – algorithmic predictions farmers cannot verify through traditional methods, creating uncertainty about when to trust technology. The dairy context exhibits amplified threat bias, where farmers are more sensitive to technological threats than opportunities. Farmers operate in an 'invisible' occupational space where the public lacks understanding of their skilled, essential work, yet their decisions carry life-or-death consequences for animals and business viability. Despite requiring sophisticated knowledge of biology, nutrition, genetics, and business management, farmers encounter stereotypes characterizing their work as 'low skill' or uneducated. These challenges create heightened threat sensitivity beyond natural human tendencies toward negative cues (Spencer et al., 2016). When technologies specifically promise to automate core expertise, threat implications become especially salient. For experts already navigating status anxieties, technologies replacing judgment with algorithmic precision represent potent threats, making dairy farming ideal for examining how pre-existing vulnerabilities shape technology interpretation.

Data Collection

My analysis centers on four farms that implemented robotic milking systems, while drawing on data from all farms to establish the baseline expertise and practices that define dairy farming. The non-robotic farms proved essential for understanding what constitutes expert dairy work – the sensory skills, tacit knowledge, and decision-making processes that farmers develop over decades. This comparative perspective allowed me to identify which aspects of expertise the robotic system threatened to displace and why certain technological features might be interpreted as particularly disruptive. Before visiting each farm, the research

team conducted preliminary interviews with owners to understand farm layout, history, and rationale for implementing robotic systems, supplementing this with research on each farm's production and specializations. Following Bailey, Barley, and Leonardi's (2021) approach for studying complex, concurrent operations, we deployed teams of two to six researchers simultaneously rather than relying on single-researcher ethnography. This team approach enabled observation of multiple activities across complex farm environments (Bailey et al., 2022; Barley, 1996). At each farm, our research team spent an entire week observing all aspects of operations (farm activities are structured in week-long cycles so this time period covered all routine operations), with multiple researchers in the field simultaneously following different informants. Each researcher shadowed one to three informants during daily routines, allowing us to understand how roles interacted and how each person interpreted work from their perspective. This multi-perspective approach captured the interdependent nature of farm operations while revealing how different actors understood and engaged with technological systems. The simultaneous observation strategy proved particularly valuable for understanding how the construction of expertise was distributed across roles, specifically, how farm managers and various co-actors performed and responded to expert performances.

Following Bailey and colleagues' (2021) suggestions, our observation protocol involved detailed running notes on farmers' interactions with technologies, people, cows, and documents. I audio-recorded technical discussions and exchanges between informants. When informants used devices (e.g., computers, PDAs, phones, tablets, robot interfaces), I photographed software interfaces, digital reports, databases, health records, alerts, data visualizations, apps, emails, texts, and spreadsheets. I collected physical artifacts including protocols, SOPs, manuals, drawings, product labels, maps, handwritten lists, forms, and

scribbled notes. Researchers kept shorthand notes marking important insights and nonverbal behaviors like toggling attention between devices and cows, allowing us to associate verbal descriptions with observed activities and assign informant-reported meaning to actions. I captured videos of complex work processes documenting tools and techniques. Since expert work was largely cognitive and internal, I complemented observation with situated questions to elucidate thought processes, interpretations of algorithmic predictions, and decision-making rationales.

To achieve 'thick description' (Geertz, 1973), I transformed observational data into comprehensive field notes through a rigorous process. At the end of each day, the research team would convene to discuss the day's activities, comparing observations of tools, activities, and interactions to create a rich contextual understanding. Researchers then expanded their running notes into detailed narratives comprehensible to someone who had not been on-site. This process involved transcribing audio files directly into the narrative at the point where interactions occurred, indexing photos of activities and digital interfaces where they were used, and weaving together actions, conversations, and images to capture farmers' gestures, movements, and the referents of dialogue (e.g., "This cow is normal…" or "Look at this cow…" [pointing to computer monitor]). This meticulous approach allowed us to document not only what farmers did and said, but also with whom, where, when, and how they performed their expertise through interactions with technology and other farm actors. Given the level of granularity required (noting minute actions down to each keystroke, text message, and verbal exchange) and the dense, complex, and interdependent nature of farm activities, completing comprehensive field notes for each hour of observation required approximately ten hours of development work.

We followed these observations with eleven semi-structured interviews totaling twenty hours, with each interview lasting between one and two hours, involving farmers, farm owners, and vendors. These interviews allowed us to question informants about specific tools, processes, or decisions in context. Interview topics included the robotic system implementation process, interpretation of algorithmic outputs, interactions with colleagues or consultants, and specific questions about tool use to understand how informants characterized their relationship with the technology. Interviews were recorded with informants' consent and had to be transcribed by hand due to the complex nature of the sound environment (e.g., cows 'moo'-ing over multiple concurrent conversations, scanners, air compressors, feed trucks, etc.), dynamic interactions, fast-paced work, and context-heavy discursive exchanges. I also gathered extensive supplementary materials including media articles about the farms and technological systems, technical research papers and documentation from vendors, industry reports and published interviews, blog posts, and technical specifications explaining the tools' features and outputs. These materials provided important context for understanding the intended functionality of the technology and industry perspectives on its implementation.

Data Analysis

Whereas the entire research team was engaged in data collection, I performed data analysis for this study. My analytical process was iterative and inductive, beginning with open coding to capture a broad range of activities, interactions, attitudes, and practices across our field sites (Glaser & Strauss, 1967). I focused particularly on interactions between farmers, technology, and other individuals on the farm, coding for instances where expertise was being performed, challenged, or reinforced.

During early stages of analysis, I was struck by the stark contrast in how farm managers engaged with similar technological systems. Although all managers were actively using the technology – interrogating predictions, incorporating it into daily routines, and integrating it into decision making – we observed marked differences in how they interpreted the technology's role in their expert work. These differences prompted me to dig deeper into the factors that might explain these divergent patterns. Our first significant analytical insight emerged when I noticed different patterns of interaction between farm managers and their various co-actors. Sometimes, others on the farm deferred to the farm managers' interpretations of cow health, and recommendations for action, while other times their recommendations were heavily scrutinized by others with whom they worked on a regular basis. While all four farm managers individually boasted over twenty years of dairy experience, it became clear that not all of them shared the same level of expert standing with their respective audiences.

This insight prompted me to examine farm managers' accounts of their relationships with co-actors more closely, revealing formative dynamics that preceded the technology's arrival. I shifted the analytical focus to these relationships, coding interactions with subordinates, superiors, and external contractors using codes such as "scrutiny," "deference," "questioning authority," "disregarding expertise," and "expressing doubt about efficacy." This analysis revealed an unexpected pattern: three of four farm managers had experienced significant challenges to their expertise from interpersonal sources unrelated to technology. Rather than assuming universally recognized expertise, I found each manager faced different threat levels prior to implementation. I returned to the data focusing on how pre-existing relational dynamics shaped technology interpretation and use. Through successive coding rounds, I identified four distinct threat patterns along two dimensions: intensity and specificity.

The DF6 farm manager enjoyed relative stability, with all co-workers deferring confidently to his expertise (**Figure 4.2a**). Managers at the other farms, however, faced threats from interpersonal sources – doubts from farm owners (**Figure 4.2b**), contempt from external experts (**Figure 4.2c**), or disrespect from employees (**Figure 4.2d**) – entirely decoupled from the technology. These varying pre-existing threats led managers to interpret and incorporate the same technology into their expertise performances in different ways.

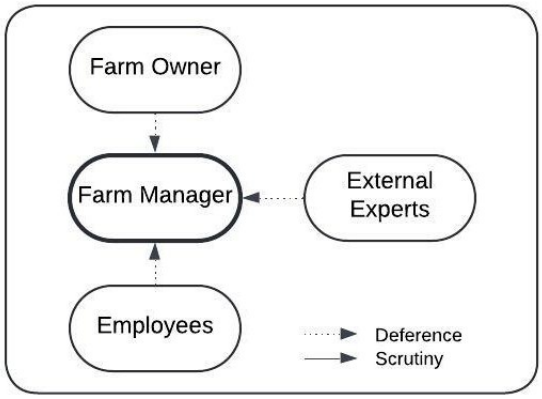


Figure 4.2a. Expertise relations at DF6.

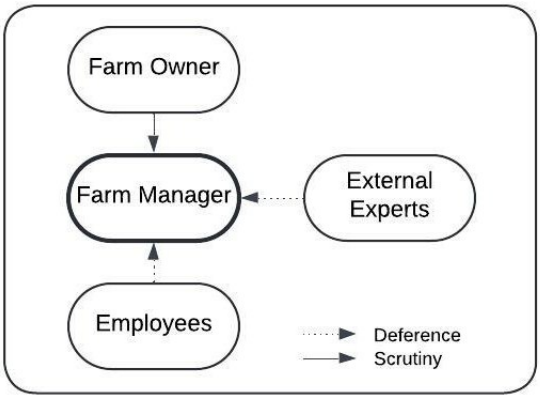


Figure 4.2b. Expertise relations at DF11.

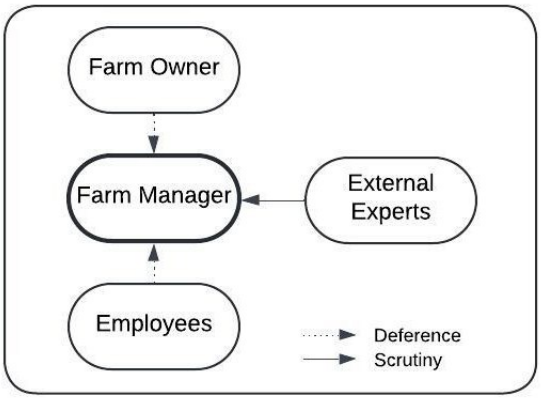


Figure 4.2c. Expertise relations at DF9.

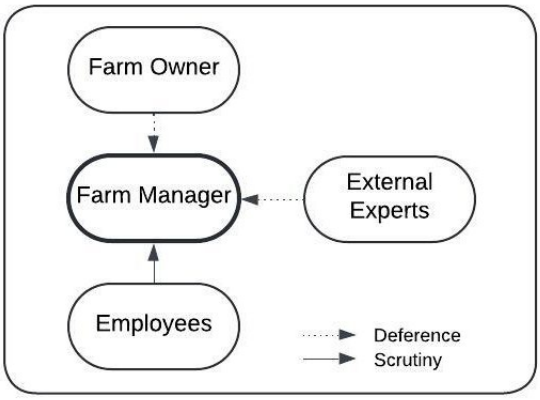


Figure 4.2d. Expertise relations at DF2.

Figure 4.2. Relational diagrams at four farms.

Through multiple iterations between the data and relevant literature, I created a characterization of threats along the intensity-specificity dimensions based on what aspects of

issues the literature would suggest contribute to appraisals of threat (Jackson & Dutton, 1988; Lazarus & Folkman, 1984). This schema showed how each farmer experienced challenges to different aspects of their expertise from non-technological sources (see **Table 4.3**). I arrived at these two key dimensions by aggregating first-order codes (e.g., "expressed doubt from farm owners," "fight with farmhand," "insulted by technician") into second-order themes (e.g., "ambiguous threat to efficacy," "intense insecurity over authority"). In doing so, I characterized threat experiences as being either high or low intensity (e.g., perceived imminence and consequentiality of threat) and either high or low specificity (e.g., identifiability and addressability of threat).

Table 4.3. Characterization of threat and how the character of threat shapes technology interpretation and use.

Character of existing threat experience at time of technology introduction	**Low intensity** *Threat is less imminent, stakes are low, and negative consequences are less existential*	**High intensity** *Threat is more imminent, stakes are high, and negative consequences are more existential*
Low specificity *Threat is distributed or unidentifiable, less salient, and less actionable because source is difficult to locate*	Daniel – Experienced little to no threat at the time of implementation. In the absence of salient threats, he oriented towards the technology as a threat and engaged in practices aimed at mitigating the threat from the technology.	Preston – Experienced intense, but unspecific threat at the time of implementation. Unsure of what to do to mitigate this ambiguous threat, he oriented towards the technology as a pathway to extend his expertise and engaged in practices aimed at mitigating feelings of anxiety and enhance his expertise.
High Specificity *Threat is narrow and identifiable, more salient, and more actionable because source is known and easier to target*	Chris – Experienced very specific, low intensity threat at the time of implementation. Hoping to combat threats from external experts, he oriented towards the technology as a symbolic resource for status enhancement and engaged in practices that recruited the technology's prestige to silence critics.	Fiona – Experienced intense, specific threats at the time of implementation. In attempts to centralize authority, she oriented towards the technology as a mechanism for control and engaged in practices that limited access to data and decision-making resources in order to maintain authority over employees.

Based on analysis, farm managers oriented toward the technology based on how they perceived its relevance to expertise. Orientations emerged as distinctive relationships between managers and the technology, influenced by pre-existing threat sensitivities and manifested in use practices. Crucially, these orientations were not predetermined by the technology's design but emerged through managers' responses to existing threats when the technology was introduced. These technological orientations influenced use practices and led to different expertise outcomes. This process revealed how managers interpreted technology through the lens of existing threats – deviating from literature that typically describes experts responding to threats from technological change itself. Instead, I found technology could be enlisted as a resource to mitigate pre-existing threats rather than simply representing a threat. This reframes our understanding of the technology-expertise relationship: rather than technology disrupting stable expertise, experts actively incorporate technology into ongoing performances to address challenges they were already navigating.

To ensure analytical rigor, we engaged in several validation practices. I conducted negative case analysis by searching for instances that challenged this emerging theoretical framework, which led me to refine my understanding of how threats influence technology interpretation and use. I also engaged in peer debriefing sessions with colleagues not involved in the study but who were familiar with the field sites to receive critical feedback or verification on my interpretations. Finally, I returned to the literature on expertise and technological change to position these findings within broader theoretical conversations, highlighting how these inductive insights extend existing understandings of the relationship between expertise and technology.

Findings

Our analysis revealed a surprising finding. Despite the predictions based on existing theory that farmers would likely interpret the technology as a threat, only one of the four farm managers, Daniel, viewed it as such. The other three, Preston , Chris, and Fiona, demonstrated strong positive orientations towards the technology and chose to engage certain features to bolster their expertise. The key difference was the presence of pre-existing threats from non-technological sources. Daniel, on the one hand, was enjoying a period of relative security in his expertise performance. He was universally regarded as an expert farm manager who was particularly effective in his role. He was frequently lauded by his boss, deferred to by consultants, and sought out for advice by colleagues. In contrast, Preston , Chris, and Fiona were each struggling with challenges to their expertise in different ways at the time of the technology's arrival. As a result, Daniel adapted his performances to combat the perceived threat that the technology posed to his expert standing, while Preston, Chris, and Fiona used the technology as a resource to mitigate threats that they were already navigating.

Although the sources of threat and patterns of behavior varied, all the farm managers went through a process of interpreting, orientating to, and engaging with the technology that were tightly coupled with the kind of threat they experienced prior to the technology's arrival (shown in **Figure 4.3**). First, prior experiences where others challenged aspects of their expertise led farm managers to interpret the new technology through what I call a "lens of threat" whereby they recognized ways in which the technology could be used to mitigate threats and reaffirm their expertise. Daniel, who was not experiencing salient threat at the time of implementation, developed a threat sensitivity to the technology, whereas the others filtered their interpretation of the technology through existing threat lenses. Second, these

interpretations led farm managers to adopt distinctive orientations toward the technology that shaped how they chose to incorporate the tool in their expertise performances. The four orientations identified in these data – critical, deliberative, deferential, and appropriative – framed which threat mitigation practices farm managers perceived would be most effective at allowing them to perform their role as experts. Finally, these technology-infused performances became self-perpetuating in the sense that the very orientations that shaped these practices reproduced not only farm managers' interpretations of the technology, but also fostered expertise outcomes that further entrenched these practices over time. I illustrate how this process played out for each farm manager in the following sections.

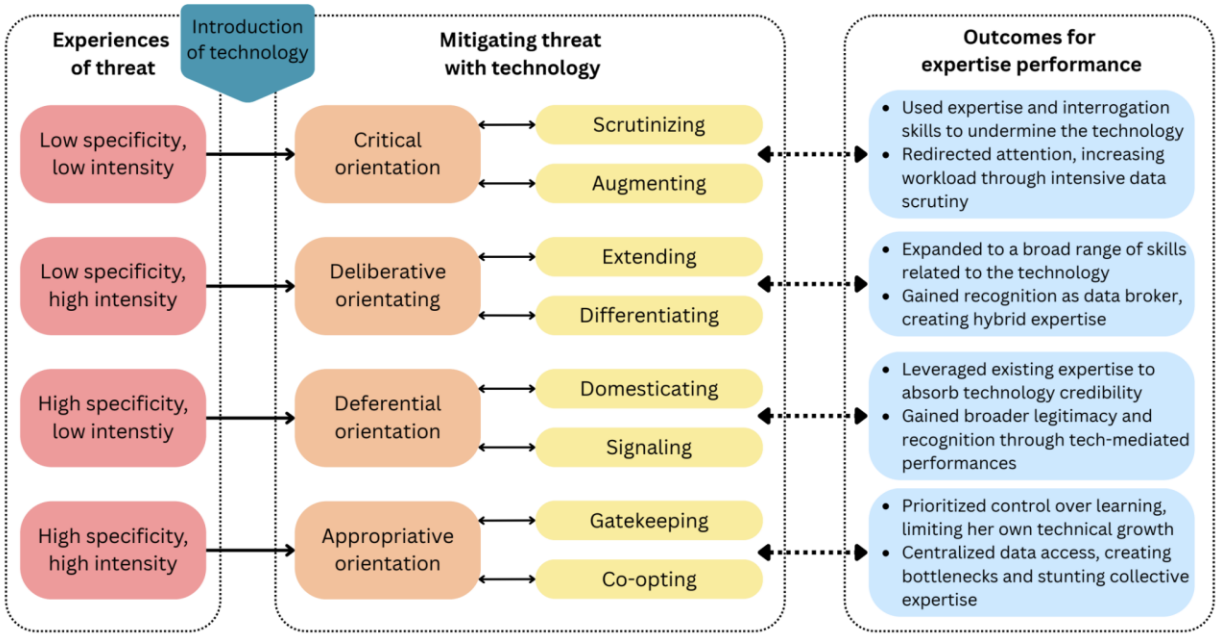


Figure 4.3. Descriptive model of threat experiences, technological orientations, threat mitigation practices, and outcomes for expertise.

DF6: Responding to Low Specificity, Low Intensity Threat

Daniel's situation was most similar to what is typically depicted in other studies of technology and expertise in that his expertise performances were relatively uncontested at the time the robotic milking system was implemented. He was widely regarded as *the* cow expert

on the farm. A humble, yet self-assured man, he described himself as having a "gift" for cows which his boss explained was what landed him the role of managing the robotic operation when the technology first arrived. Unlike the farm managers at the other farms, Daniel experienced no salient threats to his expertise prior to the system's introduction; his expertise was universally acknowledged by farm owners, colleagues, and external contractors alike who regularly deferred to his authority and recognized his superior judgment. When asked why he chose Daniel to manage the robot operation, the farm owner summed up this uncontested regard for his expertise with one simple and on-the-nose metaphor: "Daniel is just a natural cow person. He loves cows and loves to do that kind of detailed work…People rise to the top… Cream rises to the top." This unequivocal recognition for Daniel's expertise was demonstrated consistently across all members of his interdependent work network. In his case, Daniel was experiencing no explicit (unspecific and low intensity) threat at the time of the technology's arrival.

When the robotic milking system arrived, however, it directly disrupted Daniel's stable expertise performance by overlapping with the functions and recognized competencies that formed the foundation of his expert role – detecting illness and mastitis. Unlike farm managers at other farms who interpreted the technology through the lens of pre-existing threat, Daniel – who was operating in the absence of threat at the time – perceived the technology itself as a threat to the stability of his expert standing. This interpretation led Daniel to develop a critical orientation towards the technology characterized by unwavering scrutiny of the algorithm's outputs and intentional positioning of his capabilities as superior to those of the technology. While algorithmic predictions served as starting points for cow assessment, Daniel treated them as fundamentally unreliable, requiring extensive human validation and correction

which only he was uniquely qualified to conduct. He positioned himself as the technology's superior, using his decades of experience to systematically identify and highlight the system's limitations while simultaneously demonstrating his irreplaceable expertise. In doing so, two key threat mitigation practices emerged aimed at preserving his expert standing: 1) scrutinizing and 2) competing.

Scrutinizing. In an effort to maintain his supremacy over the robotic milking system, Daniel examined every algorithmic prediction systematically and with methodical precision, using multiple data sources to identify errors and applying rigorous standards that frequently led him to overturn algorithmic recommendations. He approached this practice with the intent of 'sniffing out' cases of mastitis, heat, and illness that the technology either falsely identified or missed completely. This involved interrogating the underlying data that the algorithm used to produce its claims, incorporating contextual insights the algorithm could not account for, and validating each output regardless of plausibility. This approach allowed him to demonstrate superior analytical capabilities while also exposing the technology's shortcomings. For instance, when evaluating heat detection predictions, Daniel demonstrated his detailed interrogation approach:

> "This [activity] graph helps. It gives you more data to decide [which cows are actually in heat]. You see the activity of cow #49210? It went up and it just came back down slowly compared to the activity of #54380 which just went up and back down right away. That one (#54380) is not in heat because she just went up and then right back down. This line should go up fast and then come back down slowly. But this (points to the green activity line for cow #54380) went back down really quick. She's not [ready to be bred]."

His scrutinizing practice involved examining the source data to evaluate the credibility of the algorithm's predictions. This instance revealed how Daniel transformed what could have been a simple acceptance of algorithmic recommendations into an opportunity to demonstrate his

data interpretation skills and superior judgment. Rather than trusting the algorithm, Daniel combined nuanced pattern recognition with qualitative knowledge of cow reproductive behavior to distinguish between genuine heat indicators (e.g., spike in movement because cows in heat are more active when searching for mates) and false positives (e.g., rapid spike-and-crash cycles typically seen when cows are transferred between pens which artificially increase activity recorded by the activity tracking collars). By explaining the "proper" way to read activity graphs, Daniel positioned himself as the technology's overseer, demonstrating that algorithmic outputs required his expert interpretation to be meaningful.

This scrutinizing practice extended to mastitis detection, where Daniel cross-referenced multiple data streams to identify mastitis indicators that the system missed. One case demonstrated how he developed a more sophisticated diagnostic approach that involved layering different data sources:

> "Usually when [the algorithm flags them], these numbers won't match up… Usually, they should match up like this cow. She's ten and ten (high mastitis likelihood score), but her milk is right on (points to the numbers in a data table that show average milk yield and the previous day's milk yield which are not that different). Nothing's wrong with her. That's another thing I look at."

While the algorithm produced an aggregate prediction, Daniel demonstrated analytical rigor by examining multiple metrics before accepting a mastitis diagnosis. In essence, he imposed his own system for scrutinizing and validating algorithmic predictions that he positioned as more thorough and reliable than the technology's.

Competing. While scrutinizing, Daniel also engaged in direct competition with the technology by articulating, at every opportunity, how his unique capabilities and superior judgment made him the superior instrument to detect illness, injury, and heat. He emphasized his ability to incorporate embodied, intuitive, and contextual information into his assessments

that were unavailable to the algorithm, such as knowledge of recent events that could create misleading activity patterns:

> "Those two cows went to the hoof trimmer yesterday, so they were running and then they ran back. So, even though [the algorithm] predicted that she is in heat, we need to look over here [at the event records]."

By connecting elevated activity to a visit to the hoof trimmer, Daniel demonstrated how his knowledge of farm operations and ability to contextualize data were essential for meaningful application of that data.

He often pitted himself against the technology by drawing explicit comparisons between the algorithm's and his own performance, intentionally highlighting instances where his human judgment outperformed that of the system's. Daniel asserted his dominance through these practices by routinely emphasizing the value of human intuition and embodied knowledge and experience:

> "Since I've been around animals so long, I find [mastitis] before the computer sometimes... Sometimes they're not on the heat list yet because the [activity tracker] hasn't picked it up. But I see them with my eyes... I'm out there seeing all the cows and I can see that this cow's sick... If the cow's still moving around, she's not gonna pop up on [the sick list]. I catch them earlier because I'm out there walking around."

His years of hands-on work with cows formed the bases for Daniel's competitive advantage through speed and sensitivity. Expressions like this communicate that he did not simply match the robotic milking system's performance – he exceeded it by detecting problems earlier using gut feeling and finely tuned perception. One could also detect a hint of moral superiority in his musings in how he positions the senses (e.g., "I can see that this cow's sick") and physical closeness (e.g., "I'm out there walking around") as superior instruments to devices and sensors. Daniel promotes the notion that active stewardship is preferable to passive monitoring, thereby

highlighting the irreplaceable value of embodied expertise that algorithms cannot replicate. In many cases, Daniel's competing practices extended to real-time demonstrations while walking the cow pens where he showcased how he was able to emotionally attune to the cows, a capacity the technology lacked entirely:

> "I'm just looking at their reaction. You can tell if they're sick, they'll just be standing there, their ears will be down. You know, they'll just be looking at you like, 'help me.' If they're doing well, they're just chewing their cud, just hanging out."

His ability to read subtle cues – ear position, facial expressions, body language – represented a form of expertise that remained entirely human. He often anthropomorphized the cows when talking about how he used this empathic concern and emotional commitment to their wellbeing to provide superior care.

Daniel made his competitive superiority explicit through direct dismissal of the technology's value for experienced experts such as himself. During one conversation, Daniel described the utility of the tool:

> "For people that don't know anything about animals, yeah, this is gonna help them by telling them this cow's sick. But I can find them earlier. I don't need a collar to tell me. But for people who don't know anything, yeah it'll be great for them. I'm a little more advanced so I don't need the collar."

He saw himself as someone who does not need the technology because of his extensive experience, although he does acknowledge its usefulness for those "who don't know anything," the implication being that expertise and technology use are inversely related – that real experts don't need technological crutches. This competitive framing protected his expert status by creating a hierarchy wherein technological dependence signaled inexperience.

Consequences of a Critical Orientation for Expertise

Daniel's critical orientation toward the technology produced mixed outcomes for his expertise. Through intensive scrutinizing and competing practices, he became exceptionally skilled at interpreting algorithmic outputs and detecting algorithmic errors, developing sophisticated rule-based processes for validating technological predictions. This deep engagement with the system did provide some efficiency benefits – the algorithm's lists offered useful starting points for identifying cows requiring attention, and Daniel could leverage data visualizations to support his diagnostic decisions. However, his critical approach also created a new category of expert work that involved heavy data scrutiny. Rather than simply accepting or rejecting recommendations, Daniel spent considerable time each day cross-referencing multiple data sources, applying contextual knowledge about farm events, and validating technological outputs against his observations. The technology became an additional entity requiring management and oversight, adding a significant amount of analytical work to his existing responsibilities.

This shift in attention allocation had broader implications for how Daniel performed his expertise. While he maintained his deep connection to the animals he loved and continued his regular observation routines, the added uncertainty introduced by algorithmic predictions required him to redirect some focus from more general monitoring toward intensive evaluation of individual cows. Daniel expressed particular dissatisfaction with the heat detection function, noting its limited utility, and when his attempts to collaborate with robotics consultants on threshold adjustments met resistance, he abandoned efforts to improve the system's accuracy. Instead, he relied entirely on his developed workarounds and interrogation processes, which ironically required more effort and tacit knowledge to maintain.

DF11: Responding to Low Specificity, High Intensity Threat

Although he was widely recognized as an expert with cows, Preston, the farm manager at DF11, sensed doubts from farm owners regarding his efficacy as a manager. He grew up on a dairy farm and shared stories about how sore his hands would be after helping his father milk their 50 cows as a child before continuing his dairy career at other farms. After being promoted to the manager role, he felt as though he had a lot to prove to farm owners who expressed concerns regarding his abilities to successfully manage the dairy operation. They consistently called into question his leadership and strategic capabilities. While they believed that he had potential because of his extensive experience working with cows, they privately debated his ability to apply this expert knowledge at a managerial level:

> "Preston 's perception and his focus has been pretty much staying in like, 'I know dairy, I know animals, I take care of them.' So he's really looking at it only in that lane, where we're looking at it all across [the business]. He's being more reactionary than proactive."

Here, one farm owner essentially drew a boundary around Preston's expertise, acknowledging his deep knowledge of farm operations and cows while simultaneously framing this expertise as too narrow for effective management. Her characterization of Preston as "more reactionary than proactive" suggested he merely responded to problems rather than anticipating and preventing them. The finance manager also echoed uncertainty over Preston's ability to adequately meet their expectations:

> "The guardrails for all of our other production orders are that we have these purchase orders and receipts and people who are tracking the inventory and making adjustments and managing available materials… I think all of that would be too tall an order for Preston."

This statement reinforced the threat from a different angle. While the farm owner questioned Preston's strategic thinking, the finance manager doubted his ability to even develop certain

managerial skills. For Preston, these doubts about his efficacy were not directly communicated to him, but feelings of anxiety were apparent in his daily work as pressure mounted around his ability to meet expectations that were not formally laid out to him. This meant that Preston could feel the threat to his expertise but was not able to clearly define what the threat was or how to address it.

When it came to the technology, Preston's interpretation was that it could serve as a pathway through which he could demonstrate his expertise. Because of the ambiguity of the threat, he adopted a broad approach to mitigation, which was to enhance his value in a variety of different ways. He viewed the system as a counterpart to his own expertise that both extended and highlighted his unique contributions as a skilled and knowledgeable dairyman. Like Daniel, he learned to interrogate algorithmic predictions and identified gaps and errors in the data that he then could fill. However, unlike Daniel, instead of positioning himself in competition with the technology, he combined both of their unique skill sets and knowledge bases to create a more symbiotic work arrangement. This interpretation of the technology led Preston to develop a deliberative orientation towards the technology characterized by thoughtful integration of technological insights with human judgment in order to enhance the performance of both. Preston performed the role of an expert fact checker and interpreter of algorithmic outputs, frequently combining data-driven insights with his own real-world observations. In this way, Preston demonstrated his expertise by learning to optimize and extract the maximum potential from the system. In doing so, two key practices emerged in his performance: 1) extending and 2) differentiating.

Extending. In an effort to enhance his expertise performance, Preston used the technology to extend his capabilities. Preston spent hours each day interpreting algorithmic

predictions to curate task lists, learning the technical aspects of the software to better interrogate its outputs, and exploring different visualization capabilities to gain new data-driven insights about cow health trends. He extended his expertise by carefully studying the tool's outputs, identifying patterns that indicated false positive predictions and others that allowed him to identify individual cows that needed medical attention. He sought out opportunities to learn how to repair the robots to extend his value: "I asked the [electrician] what I have to do to fix the sensor [in the robot]... Sometimes they know but don't want to tell me so I keep pressing! Otherwise, I just figure it out myself." Similarly, he always wanted to learn new capabilities in the software program: "I asked [the software consultant] to show me how to make my own reports in the program because I want to know how to make them myself." This proactive and deliberate pursuit of new knowledge and skills allowed him to extend his expertise by performing new data analytical, technical, and mechanical expertise to farm owners. In some cases, he leveraged data from the system and combined them with his own tacit knowledge in creative ways to find injured cows:

> "When you see high numbers like these go up [points to a column in a data table], I think something happened. Maybe one robot is not milking well and you can start checking the robots or looking at the cows because sometimes the same cows have problems. So, you can look to see if it's the robot or if it's the cow. Sometimes when some cows go lay down in the bed and another cow goes to lay down [next to them] and that other cow's udders are really big. The cow goes to lay down and they accidentally step on the [udders of the other cow]. Then that cow can't milk any more [because her teat is injured]. Then the robot starts catching this cow as a failure because they don't get any milk out of that teat. We need to give something for pain or medicine for infection… We can catch that by looking at the [robot] failures."

This example showed how Preston extended his expertise by creatively repurposing the technology's data. He recognized that data about milking failures could be used to identify udder injuries, thereby extracting unexpected insights from the data. Most notably, Preston

combined his tacit knowledge of cow behavior (understanding the uncommon injury mechanism where one cow steps on another's udder) with data (robot failure reports) to create a novel diagnostic capability that neither the technology alone nor traditional observation could achieve. Rather than discounting the data, Preston actively extended his capabilities by integrating technological insights with his specialized knowledge in ways that enhanced his expertise.

Differentiating. While extending his capabilities through the technology, Preston also made sure to differentiate aspects of his expertise that the technology could not replicate that complemented and enhanced its effectiveness. He demonstrated his specialized knowledge by carefully evaluating the algorithm's predictions and applying context-based judgment. Preston exemplified this approach during one observation session when reviewing the system's heat predictions:

> He clicks on one of the lists which shows all the cows that are likely to be in heat – there are 6 cows on the list today. He points out the heat probability score for cow #1876 – a 96 probability score. The other scores are a lot lower and he says that these cows' activity levels are probably slightly elevated. "Maybe the cow ran a little bit and the robot said maybe it is in heat," he says. But he tells me they are likely not in heat. [Excerpt from field notes]

This daily interrogation practice differentiated his expertise by applying critical judgment to the algorithm's predictions. Like Daniel, he rarely accepted the algorithm's predictions at face value. But unlike Daniel, he did not do this in order to compete with the technology or undermine its value. Rather, he appreciated the affordances of the technology while maintaining his optimal distinctiveness as an essential interpreter whose specialized knowledge complemented and enhanced the technology rather than being devalued by it. Preston saw his knowledge and the technology's data as complements. When explaining his ability to detect illness visually – "When you look at the cow, you can tell. They look sad, they

look tired, they can hardly move. The udder is big in one quarter and sometimes they look blue" – he was not merely describing traditional skills but asserting his specialized perceptual expertise that algorithms could not replicate while incorporating data to verify his intuition. He also differentiated himself through personalizing his relationship with the cows. He often emphasized being able to recognize individual cows as a unique ability of his: "Maybe to you, all the cows look the same. To me, every single cow looks different." Further, when he identified that a farm-wide milk production decline was caused by a feed composition change rather than illness as the algorithm had suggested – "The cows really like the brewers grain so when we aren't able to add it, they don't eat as much and this is reflected in milk production" – Preston demonstrated the systems-level thinking and proactive analysis that management had questioned. Preston categorized the technology as an opportunity to extend and differentiate his expertise to combat this ambiguous and high intensity threat by incorporating more general and pragmatic practices into his performance.

Consequences of a Deliberative Orientation for Expertise

Preston 's deliberative orientation towards the technology transformed his expertise in ways that enhanced his capabilities. With access to more granular data on individual cows through the technology, Preston developed even more detailed knowledge of individual cows and their specific idiosyncrasies, as well as new skills working with data and robots. He explained, "When you see high numbers like if I saw these numbers go up, I think something happened [either with the cow or with the robot]," demonstrating how he began connecting digital indicators with physical realities, expanding his interpretive capabilities and decision-making framework. His relationship with the cows became more mediated through data, yet paradoxically more individualized. Through this continual validation and contextualization of

algorithmic claims while in the barn, Preston developed sophisticated interpretive skills that complemented rather than replaced his traditional expertise.

As Preston's expertise performance became more analysis-based, he became known on the farm as the data expert. By becoming proficient at validating and contextualizing technological outputs, he established himself as an essential translator and data broker between different departments on the farm. His role expanded beyond traditional animal husbandry to include aspects of data interpretation and communication, creating a hybrid expertise that was uniquely valuable to the organization. As a result, Preston's approach resulted in genuine expertise expansion. By treating the technology as a complementary knowledge source rather than a competitor or replacement, he developed a more sophisticated, hybrid expertise that combined technological literacy with deepened domain knowledge.

DF9: Responding to High Specificity, Low Intensity Threat

While Preston experienced doubts from within his farm, Chris, the farm manager at DF9, faced challenges from outside the farm. His work often required collaborating with external consultants, contractors, and vendors to resolve problems with equipment or make strategic decisions about the farm. Described as a "micro manager," Chris was proud of his identity as an expert dairyman and made an effort to be highly involved in day-to-day work on the farm. He had decades of experience working with cows that earned him a management position at the farm. His employees and boss all regarded him as an expert in farm related matters, but not everyone shared this attitude outside the farm. Some contractors did not view Chris's status as an expert dairy farmer as professional or legitimate, making cutting remarks that belittled his expertise. This status devaluation created a persistent threat to his legitimacy, as revealed in his frustrated description of interactions with an IT consultant:

Chris: There are some [consultants] that are really stuck up, you know like they think, 'oh, these [farmers] don't know what they're doing.' They look at you like you're dumb.

Researcher: You've been working with cows for over 30 years, but do they not count that as expertise?

Chris: No, they don't.

He threw his hands up in indignation conveying both the clarity and the emotional weight of this dismissal. This was not a unique interaction.

Over the years, Chris felt that these external contractors looked down on his knowledge and cow experience because he did not have a formal education. "[This software consultant] has his fancy degree. I didn't go to school but I have a good job, I'm paid well, I have a good boss, and good job satisfaction," he explained in defense of his experience. This statement reflected his awareness of how educational credentials were being used to undermine his professional standing despite his successful career and practical accomplishments. Even material symbols of his success became battlegrounds for status, as when a pharmaceutical representative challenged his worthiness of a new truck gifted to him as recognition from the farm owner – "You shouldn't have this new truck. [The farm owner] should have that truck." This remark, implying that Chris's position did not merit such recognition, represented a direct attack on his expert standing. Despite being recognized as an expert on the farm, he occasionally had to contend with threats from outsiders – an annoyance that was not overtly damaging to his expert status but was pernicious enough to affect his sense of self, nonetheless. In this sense, the threat he experienced was relatively diffuse, but stemmed from a specific and identifiable source: outside experts.

When the technology arrived, Chris was already contending with these threats to his expertise and quickly interpreted the technology as a mechanism to prove his expertise to

naysayers. Where these professionals with "fancy degrees" had dismissed his practical knowledge as inferior, Chris saw an opportunity to redefine his expertise around mastery of the very technology these consultants claimed as their domain. Rather than viewing the technology as something that might threaten his value or relevance, Chris recognized the symbolic prestige of mastering such an advanced technology that even external experts could not easily dismiss. This interpretation of the technology as a proving ground led Chris to develop a deferential orientation characterized by enthusiastic public endorsement of the technology's sophistication while simultaneously positioning himself as uniquely qualified to optimize its performance. In doing so, two key practices emerged around how he engaged with the technology: 1) domesticating and 2) signaling.

Domesticating. In domesticating practices, Chris actively demonstrated his mastery over the technology by applying organizational knowledge to adapt the system to the specific context of his farm. When provided with generic implementation protocols from the robot manufacturer and their technology consultants, he rejected their approach:

> "Well, [the robot consultants], they have their ideas. And I did not agree with them... So, I came here to the office and I said to [the farm owner], 'I don't want nobody from outside. Give me a week to train the cows.' He's like, 'Are you sure?' And I said, '100%... So, he calls the [robot consultants] and he tells them that's what we're going to do in the next barn. So we don't need nobody from outside. 'We don't need you guys here.' And they're like, 'Oh, he's crazy. It's not gonna happen.' So [the farm owner] calls me, 'Hey, they told me you're crazy. It's not gonna happen. Just watch if it's gonna happen or not!'"

Chris took control of the implementation process from here because he believed that external experts did not possess the necessary expertise to successfully transition their cow herd – a process that requires round-the-clock monitoring and training to acclimate the cows to the milking robots. By directly excluding outsiders who he viewed as a threat, Chris was actively

reclaiming his authority. He wanted to make a statement about whose expertise truly mattered in this context – his deep, tacit knowledge of their farm and cow behavior, not abstract technical training – and his approach proved remarkably successful:

> "We started Monday. By Wednesday, three days later, I'm pulling the gates out because the cows already go [in the robots] on their own. They said we had to run them for two weeks... In three days, I was done!"

His successful transition program not only demonstrated a deep understanding of how to train cows, but also his ability to domesticate the technology which was a feat that the robotic consultants did not believe was possible. This achievement allowed him to reframe the relationship between his traditional expertise and the technology, positioning himself as uniquely capable of bridging these domains in ways that challenged the presumed superiority of technical specialists. The dramatic outperformance of the consultants' timeline – completing in three days what they said would take two weeks – served as irrefutable evidence that his hybrid expertise was not only legitimate but superior. For Chris, who had been made to feel that his three decades of experience "didn't count as expertise," this concrete achievement represented a powerful status reversal, transforming him from someone perceived as an uneducated farmer to someone whose expertise was demonstrably more effective than that of formally educated specialists.

Signaling. Simultaneously, Chris engaged in signaling practices, strategically communicating his expertise to relevant stakeholders. He raved about the technology – "With the robot, there's no guessing. So, [the breeder] is not picking. The robot is telling him, 'You breed these cows.' That's it. Don't touch nothing else. It's amazing!" These statements served not merely as praise for the technology, but as strategic declarations that elevated his own status by association with advanced technological systems. He displayed his mastery over the

technology by citing performance metrics and promoting the system's capabilities: "If [the robot] tells you there's mastitis in two quarters, she has mastitis in two quarters. It's amazing!" This endorsement directly countered the threat he experienced when external contractors viewed him as uneducated with rudimentary skills. He positioned himself as more advanced than peers who resisted similar technologies, noting, "A lot of people don't want to use [the algorithm]. They won't breed out of the robots... Why? Our conception [rate] is high now since we went to robots. About 33%!" These signals served a dual purpose – promoting the technology while simultaneously signaling his superior expertise in association with successfully mastering it.

At the time of implementation, Chris was experiencing a very specific threat to his expertise performance stemming from external consultants, but on the farm, he was still hailed as an expert. This specific, low intensity threat shaped his mitigation strategies to be more directed at those consultants who had disregarded his expertise and the way he chose to mitigate those threats was to neutralize the primary critique they imposed about the value of his cow expertise. In this case, the technology became integrated into Chris's new hybrid role in which the tool became a critical component of expertise performance, allowing him to project an image of a sophisticated "robot farmer" to external actors. The transformation from traditional dairyman to robot farmer offered a powerful counter narrative to the derogation he had experienced when contractors implied his work was unprofessional or illegitimate.

Consequences of a Deferential Orientation for Expertise

Chris's deferential orientation positioned the technology as a symbol of expertise that enhanced status and legitimacy. He managed to combine technological expertise with cow knowledge, creating a new professional identity as a "robot farmer." His approach involved

developing substantive understanding of how the technology functioned in practice insofar as to be able to effectively perform mastery of it to audiences whom he thought might dismiss his expertise. This hybrid identity dramatically reconfigured Chris's work relationships, particularly with external consultants and technicians. As he explained regarding his relationship with a software consultant:

> "He knows the computers and I know the cows. We need each other. He just acts like he knows everything... It didn't take long [to convince him]... I wanted to do it one way and he wanted to do it a different way. But I'm here with the cows every day so I know what they need... I think [the consultant] learned his lesson."

This exchange illustrates how Chris was able to convince his critics that his cow expertise was critical to the success of the technology in such a way that enabled mutual recognition. Rather than categorizing the technology as a threat, he performed an essential role by asserting the very expertise that external consultants did not recognize as legitimate by foregrounding it in the context of the technology, which paradoxically earned him recognition as a robot expert.

Through these performances, Chris cultivated wider professional recognition that transcended the farm boundaries. His successful domestication of the technology became tangible evidence of his innovative approach and gained formal recognition when he earned the "Farmer of the Year" Award at an international dairy technology convention, expanding his expert reach beyond local contexts as other farmers in the region also looking to implement robots sought out his expertise as a consultant. Chris's performance incorporated the technology as an institutional signal of expertise which resulted in an emergent role that created new possibilities for expert recognition in modern dairy farming.

DF2: Responding to High Specificity, High Intensity Threat

The last farm case takes place at DF2, where the farm manager/owner Fiona had a long history of herdsmen disregarding her expertise. For Fiona, this threat was both incredibly intense and highly specific in that she knew exactly what/who was threatening her. As a woman in a predominantly male industry with a culture of "machismo," Fiona struggled to achieve the recognition that she believed she had earned in her decades of committed work on her family farm. Despite her father wanting to leave the farm to her male cousin, she "worked her way up" and earned a management role before assuming ownership of the farm when her father retired years prior. Over the course of her tenure, she had several altercations with older dairymen who did not respect her expertise and lashed out in aggressive displays of insubordination. The impression that these incidents had on her was apparent in how she spoke about the interaction, even years after the other person resigned. She described one such incident with a former herdsman:

> "[This herdsman] was authoritative and aggressive. He used to tell me that I spent too much time in the office. I was raising three kids and running the dairy and I told him it was not for him to tell me what to do… He was very threatening. Everyone is a lot happier now [after he left]. He had a machismo attitude. The [other herdsmen] told me afterwards that he always did the opposite of whatever I told him to do."

This criticism stuck with Fiona as she shared other stories about similar confrontations where herdsmen would yell in her face. These experiences of explicit disrespect created a feeling of persistent threat that left Fiona suspicious of her own employees. Even after years of managing the farm, Fiona was still sensitive about these experiences and tried to maintain control over decision making on the farm to cope with feelings of deep-seated insecurity. For Fiona, these salient challenges to her authority led her to become a self-proclaimed "control freak," explaining, "I need to review everything first," as she looked for ways to preserve her expert

authority. Dealing with profound threats to her authority had made Fiona very protective of her ego to the point where she did not trust others even to translate her instructions into Spanish:

> "My Spanish is terrible, but I can generally follow. I can translate with my phone. I ask [the herdsmen], 'Are you sure you understand what I'm saying? You have to tell me.' I try to find a way to say it so they understand. No handholding, and I don't trust people translating for me."

Even when it came to delegating tasks to her employees, she wanted to be in full control over communication and confirmation to ensure compliance among her workers, demonstrating how thoroughly her previous experiences of insubordination had shaped her management approach, leading her to prioritize control over delegation even in basic communications.

When the technology arrived, Fiona was still reeling from the long-term impacts of threat and because her defensive attention was focused on maintaining control over her employees, she interpreted the technology as a tool to maintain control over all farm decisions, viewing it as a resource that could reinforce her expert position rather than as a threat itself. Fiona identified a critical governance role that she could perform in curating lists and evaluating algorithms outputs by centralizing information flow and positioning herself as the essential intermediary between the system and her employees. This interpretation of the technology led Fiona to develop an appropriative orientation characterized by instrumental control where she limited others' ability to engage directly with the system by restricting access to data and selectively leveraging the technology's veil of objectivity. By employing these features, two key practices emerged around how she engaged with the technology: 1) gatekeeping and 2) co-opting.

Gatekeeping. In gatekeeping practices, Fiona actively restricted access to the system's data and outputs, positioning herself as an essential intermediary between the technology and farmhands. Despite acknowledging her limited technical background, she strategically

developed just enough familiarity with the system to maintain control over its outputs. Her admission – "I never learned Excel... I am a bit of a control freak; I like to do everything myself" – revealed how her technological approach directly stemmed from her authority concerns rather than from interest in the technology itself. The self-identification as a control freak particularly illuminated how she viewed centralized information control as essential to maintaining her position. Rather than providing workers direct access to the system's outputs, Fiona deliberately created an information bottleneck by personally curating all data before distribution. She consolidated information across different pens and distributed customized lists via text message, manufacturing dependence on her interpretation of the data. This practice transformed what could have been a democratizing technology into a tool that reinforced hierarchical relationships. The effectiveness of this strategy was evident in one herdsman's explanation:

> "We need to wait for Fiona to send them a fetch list she creates. He says that he can also see the list of the cows to be collected by generating a report through the robot's monitor... But this is the list that Fiona is currently looking at to create her own fetch list."

This statement from a herdsman revealed the artificial nature of workers' dependence – they technically could access the information themselves but had been conditioned to wait for Fiona's curated lists. By gatekeeping workers from developing their own skills with the technology, Fiona effectively addressed her authority concerns by eliminating potential challenges to her interpretations and ensuring her employees had no choice but to defer to her expert judgment.

Co-opting. Simultaneously, Fiona strategically appropriated the technology's capabilities to reinforce her authority through surveillance and rationalization. Even though she had a limited understanding of how the algorithm actually produced its predictions, she

leveraged the technology's perceived objectivity as a shield against challenges to her decision to defer to it. When the breeder questioned the heat detection predictions, she responded by invoking the system's technical documentation rather than engaging in substantive discussion about its accuracy:

> "I asked about the heat score because a breeder wanted to know how they calculated it because he disagreed with the values... I showed the breeder the sheet with information explaining the algorithm's logic."

Despite not understanding the algorithm's inner workings herself, she used its technical documentation as a proxy for expertise, effectively deflecting the breeder's challenge by replacing her contested personal judgment with the seemingly objective authority of the technology. This approach allowed her to maintain her position even when facing disagreement from other experts who might otherwise dispute her.

Fiona also repurposed the technology's features to address her longstanding concerns about worker compliance. By examining the "time away" reports in the milk tracking software which monitored how long each cow was away from the milking robot, she transformed a milk production monitoring metric into a worker surveillance tool:

> She is looking at the 'time away' column and is wondering whether the night crew fetched the cows they were supposed to... Now she compares the daily report to yesterday's fetch list to see if the night shift actually fetched the cows (resulting in these high numbers if they did not, suggesting the collect cows list serves as a management surveillance check as well). She says that to check her suspicion she is going to look at last night's [handwritten] fetch list to see if the same cows appeared there. [Excerpt from field notes]

This creative repurposing of data allowed her to passively monitor compliance remotely rather than through direct confrontation and oversight.

At the time of implementation, Fiona was still experiencing the lingering sting of very extreme and pointed criticism from her employees. For her, the threat still felt very intense and

very specific to her interactions with her subordinates. As a result, she appropriated specific features of the technology that allowed her to aggressively target the source of threat by centralizing authority, restricting employees' access to data, and surveilling their work. Through these practices, the technology became not just a tool but a substitute for Fiona's expertise, allowing her to perform an authoritative role while limiting the need to engage directly with workers.

Consequences of an Appropriative Orientation for Expertise

Fiona's appropriative orientation produced mixed outcomes for her expertise. By positioning herself as the essential intermediary for all technological outputs in an attempt to leverage the system's rational authority and surveillance capabilities, Fiona effectively reconfigured her expertise performance to prioritize control over knowledge enhancement. As a result, she developed minimal technical proficiency while becoming increasingly detached from the practical work in the barn that had initially formed her expertise base. With increased focus on managing information flow through the software system, Fiona spent more time in the office, ironically validating past criticisms that had threatened her in the first place (i.e., the previous incident where a herdsman told her, "You spend too much time in the office."). Fiona's selective engagement with the technology prioritized control functions rather than developing a deeper understanding of the system's capabilities to enhance her ability to manage the farm. When questioned about how the algorithm produced its outputs, Fiona demonstrated limited technical knowledge:

> When I ask why the heat probability threshold is set for 25 (especially when they look for numbers much higher), she tells me, "I'd have to ask her [the robot rep]." She clicks around on the screen looking for an answer for me but finds nothing useful. [Excerpt from field notes]

This superficial technical understanding meant her expertise became increasingly dependent on the technology's outputs without the corresponding ability to evaluate or contextualize those outputs independently. This gatekeeping approach extended to her relationship with external data sources, as she eliminated alternative information channels that might provide competing perspectives:

> She had been paying $2000 monthly to belong to the Dairy Herd Improvement Association, which published production numbers and gave out awards and conducted a testing program. She did her own test and found out that they were making mistakes, so she stopped... "My robot is going to tell me everything that I need to know." [Excerpt from field notes]

This consolidation of information sources narrowed the knowledge ecosystem informing her decisions, limiting her exposure to alternative approaches or industry developments.

Fiona's practices reshaped her relationships with workers in complex ways. Appropriating he surveillance affordances of the technology enabled oversight but created distance from the work and people in her expertise network. By positioning herself as an intermediary between technology and farmhands, she created dependencies that upended work structures. Workers became reliant on her curated lists rather than developing their own skills with the technology, creating an expertise bottleneck where all technological insights flowed through her. This centralization constrained the farm's adaptability and collective learning, as workers remained disconnected from the technology's capabilities. While her position became more dominant, employees' roles blurred. She put younger, inexperienced workers in charge of robotic operations – choosing those unlikely to push back – which confused both these workers and senior dairymen about their relative positions as less experienced workers took on seemingly advanced tasks around robot maintenance. Through performances aimed at alleviating specific threats, Fiona inadvertently prioritized control over understanding. Her

gatekeeping and co-opting practices positioned technology as a substitute that both reinforced and confounded hierarchical relationships while stunting practical expertise development. Farm workers couldn't develop data and analytical skills, while Fiona became so overburdened as sole interpreter of algorithmic outputs alongside managerial responsibilities that she lacked time to advance her own expertise. Rather than enhancing expertise, her focus on control led to maladaptive coping strategies that limited her ability to develop the competencies that would earn the recognition she desired. The technology became a tool for maintaining threatened authority rather than developing new capabilities.

Discussion

This ethnographic investigation of dairy farms adopting robotic milking systems revealed that farm managers experiencing ongoing threats when technology arrived incorporated the tool into their performances to mitigate these threats. Surprisingly, despite the technology exhibiting all features expected to trigger threat categorization – high potential for negative consequences, low control and competence feelings, and ambiguous meaning (Jackson & Dutton, 1988) – three of four farm managers interpreted the technology positively rather than as threatening. Only one demonstrated the typical threat response to the technology itself. The other three interpreted the technology as an opportunity and enlisted it as a resource to combat existing threats in ways aligned with the specific threat type they sought to mitigate.

Figure 4.4 depicts the theoretical model showing how existing expertise performances and continuously navigated threats shape technology interpretation through what I call a *lens of threat*. Past threat experiences create sensitivities and attentional biases that filter how experts categorize technological threats and opportunities. In this study, farmers altered their performances to incorporate technology differently, producing significantly different expertise

consequences. These patterns reproduced threat mitigation practices as they became entangled in expertise performances, reinforcing the technological orientations that produced them. These findings demonstrate how technology can be strategically enlisted to address different threat types, and how introduction timing relative to threat experiences has important implications for expertise performance and technology use. The resulting picture of the technology-expertise relationship proves far more nuanced than current theory acknowledges in that experts' interpretations are profoundly shaped by ongoing performances beginning long before technology arrives. I outline the implications of this study for two key areas of focus for scholars of technology and expertise: technology interpretation and outcomes for expertise.

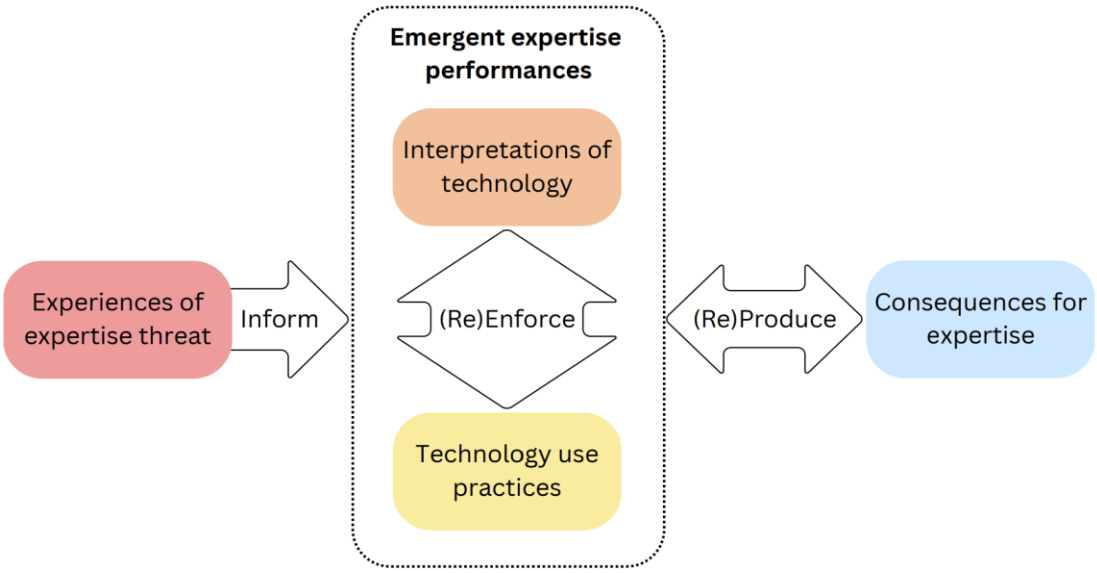


Figure 4.4. Theoretical model of compensatory technology adoption in response to expertise threat.

Interpreting Technology through a Lens of Threat

Prior research has repeatedly argued that new technologies threaten expert standing across diverse contexts and occupations (Anthony, 2021; Beane, 2019; Bechky, 2020; Nelson & Irwin, 2014). This conclusion is reasonable given that technologies meet the established

criteria for stimuli that trigger threat categorization (Jackson & Dutton, 1988) and that modern technologies increasingly perform functions that I traditionally associate with experts, such as making complex predictions, analyzing and interpreting information, and rendering judgments that once required years of training. Thus, it is logical that the majority of studies examining technological change and expert work default to treating technology as the primary threat to which experts respond. However, expanding the temporal aperture to examine how expertise performances are already under threat before technology arrives reveals important social factors that prime experts to interpret new technologies in particular ways.

Our findings show how experts' past experiences managing threats to their expertise shape how and whether they categorize new technologies as threats or opportunities. The two dimensions of threat – intensity and specificity – are central to understanding why experts interpret the same technology in radically disparate ways. These dimensions matter because they show that it is not the technology itself that influences whether it is categorized as a threat or an opportunity, but the kinds of threats to expertise that experts are already managing when the technology arrives. High-intensity threats generate strong defensive responses because they make the stability of one's expert standing feel precarious. High-specificity threats, in turn, provide clear and identifiable targets for mitigation, shaping how experts choose to incorporate technology into their performances. Taken together, intensity and specificity provide a theoretically important framework for explaining why experts enlist technology as a resource for preserving their authority in some cases and reject it as a rival in others.

Table 4.3 illustrates how these dimensions operated across these cases. Crucially, however, intensity and specificity are not fixed attributes of an expert's standing; they fluctuate over time as expertise performances become more or less contested. At some moments, experts

perform primarily to parry salient threats, while at other times their performances unfold in relative stability. Daniel was the only manager situated in the low-intensity, low-specificity quadrant at the moment of implementation, meaning his expertise was uncontested and secure. Put another way, he was engaged in a relatively well-scripted and uncontested performance of expertise when the robotic milking system arrived. Precisely because he was not already defending against salient threats, he perceived the technology to present a novel and direct challenge to his role: it encroached on the very diagnostic practices that underpinned his uncontested standing. Unlike the other managers, who interpreted the system through the lens of ongoing struggles and used it to shore up their positions, Daniel had no such vulnerabilities to mitigate. For him, the technology itself became the threat, and he oriented himself to it with scrutiny and competition. This divergence underscores the theoretical significance of intensity and specificity: they shape not only the interpretation of new technologies but also why only certain experts perceive them as threatening in the first place. Importantly, these findings suggest that a different technology, arriving at a different point in the flow of any of the farm managers' expertise performances, could have been interpreted in very different ways.

Our findings also show that social identity attributes shape experts' exposure to both the intensity and specificity of threats and how their expertise performances are perceived by others. Intense threats heighten defensiveness, while specific threats provide identifiable targets for mitigation (Jackson & Dutton, 1988; Lazarus & Folkman, 1984). But identity characteristics such as gender, education, and ethnicity moderate both the degree to which individuals are vulnerable to threat and how individuals experience particular kinds of threat – sometimes making them exacting and more personal, other times more diffuse and stigmatizing (Petriglieri, 2011; Shih et al., 1999; Spencer et al., 2016; Walton et al., 2015). For example,

Fiona experienced highly intense and specific challenges to her expertise rooted in gendered authority struggles, while Chris's expertise was discounted due to the more diffuse but specific threat of lacking formal education (Brescoll, 2011; Kraus et al., 2017; Stanaland et al., 2023). These differences shaped not only the specificity and intensity of the threats they encountered but also how audiences interpreted the credibility of their performances due to systemic inequities in how expertise is recognized in certain individuals and not others. By contrast, Daniel, whose expertise was widely recognized and uncontested, operated in a context of low-intensity, diffuse threats, which made the technology itself appear as the primary danger.

These dynamics extend prior work that treats expert responses as simple dichotomies, such as opportunity versus threat (Van Doorn et al., 2023), acceptance versus dismissal (Anthony, 2018; Glikson & Woolley, 2020; Lebovitz et al., 2022), or aversion versus appreciation (Dietvorst et al., 2015; Logg et al., 2019) – by showing that interpretations are mediated by the interaction between social identity, threat type, and audience perception of expertise performances. Experts do not engage technology based on inherent features alone but enlist it selectively to counter threats of varying intensity and specificity and to perform their expertise in ways that secure recognition from relevant audiences. This perspective explains both the heterogeneity of responses and why only some experts perceive technology as threatening. In doing so, it reframes the technology-expertise relationship from one of disruption to one of strategic incorporation shaped by the ongoing management of identity-laden threats.

These findings suggest we might reinterpret past studies of technological change and expert work to explain not only *how* but also *why* experts responded to technologies in particular ways. A threat lens could clarify differences in emergent practices that prior research

observed. For instance, Barley's (1986, 1990) classic study demonstrated how new imaging technologies transformed radiologists' and technologists' roles, restructured networks of interaction, and ultimately realigned occupational hierarchies. His account showed that new modalities redistributed interpretive expertise toward younger radiologists and created new distinctions among technologists, thereby unsettling taken-for-granted role relations. These findings suggest that these dynamics can be more precisely understood through the lens of threat intensity and specificity. The arrival of ultrasound and CT introduced threats of varying intensity and specificity to different cohorts: for younger radiologists, the threat was diffuse and low in intensity, as their training already encompassed the new modalities, whereas for medium-tenured radiologists, the threat was both intense and specific, challenging their interpretive competence in direct and visible ways. Older radiologists, by contrast, often experienced threats of low intensity, since their established status and limited expectations for retraining muted the degree of contestation they faced. Seen this way, the upheavals that Barley documented were not simply the result of role reconfiguration or network realignment, but of how technologies shifted the threat landscape – who was exposed to high-intensity, specific threats and who could rely on more stable, low-intensity positions

Rethinking the Link Between Technology Use and Expertise Performances

Prior studies show that incorporating new technologies into expert work transforms expertise itself by reshaping cognitive processes (Anthony, 2021; Lebovitz et al., 2022), interactions (Barley, 1986; Leonardi, 2007), and performance repertoires through which expertise is enacted and recognized (Leonardi & Treem, 2012; Treem, 2013). These findings extend this work by showing how experts can become unwittingly entrenched in performative roles when threat mitigation takes precedence over broader considerations of its consequences.

In adapting how they demonstrate knowledge, apply skills, and engage with tools, experts reshape both what they know and how they perform expertise (Beane & Anthony, 2023). Findings also reveal that the content of expertise can shift based on responses to prior threats: Mitigation strategies, when narrowly focused on specific challenges, can lock experts into particular performance patterns.

Figure 4.3 shows how this process unfolds. Experts' positions in the threat landscape relative to the intensity and specificity of the threats they faced can shape the orientation they adopt toward the technology. In this study, low-intensity, diffuse threats led to a critical orientation, in which experts engaged in scrutinizing or competing practices. High-specificity threats pushed experts toward deferential or appropriative orientations, in which they engaged in domesticating, signaling, gatekeeping, or co-opting. Those navigating ambiguous or intermediate threats adopted a deliberative orientation, engaging in extending or differentiating practices. These findings build on work that has shown experts use technologies to bolster or extend their standing (Prasad, 1993; Treem, 2013), but depart from accounts that portray such moves as discretionary. Instead, I show that orientations are channeled by prior threat experiences, which can constrain how technologies become incorporated into expertise performances.

What is most important about these orientations is that they do not appear to stop at threat mitigation. On the farms I studied, the practices tied to each orientation became central features of experts' ongoing performances. Scrutinizing outputs positioned some experts as guardians of rigor, making interrogation and critique the hallmark of their expertise. Signaling technological mastery allowed others to gain legitimacy but narrowed their expertise to visible displays of competence. Gatekeeping data flows reinforced authority but reshaped expertise

into control rather than knowledge creation. In each case, the technology became embedded in performance repertoires in ways that altered the very substance of what expertise looked like to audiences. Prior work has documented how technologies reconfigure expertise roles (Anthony, 2021; Barley, 1986; Nelson & Irwin, 2014), but has not fully accounted for the mechanism by which technologies reshape the content of expertise itself. This study shows that these shifts emerge not from the technology alone, but from how experts, who are situated in particular threat landscapes, mobilize technologies as resources in their performances.

Over time, these orientations produced divergent trajectories of expertise, even among experts with similar knowledge bases engaging with the same technology. Because practices such as augmenting, signaling, or gatekeeping became recognizable markers of authority, experts were typecast into specific roles. These roles both reinforced their standing and limited adaptability, making deviation risky. Thus, a critical orientation could transform an expert into a perpetual skeptic, valued for interrogation but sidelined from innovation. An appropriative orientation could bolster authority through co-optation but foreclose collaborative learning. These findings diverge from prior work that treats expert responses as binary and instead show how different threat experiences generate distinct performance paths.

Taken together, this analysis demonstrates that experts' locations in the threat landscape shape the orientations they adopt, which drive the activities they engage in with technology, and these activities feed back into performances that redefine expertise. This recursive process clarifies how the same technology can generate different outcomes for experts with similar knowledge bases. The prop they make of the technology depends on their orientation, and the performance they build around that prop reshapes their expertise itself. In doing so, I extend work on the performativity of expertise (Bourgoin et al., 2020; Eyal, 2013; Huising, 2015;

Treem, 2016) by showing that newly implemented technologies can set in motion feedback loops that transform what expertise is. Findings suggest that new technologies create divergence in expertise trajectories, as experts' positions in the threat landscape channel them into orientations that both transform and entrench their performances.

Conclusion

The farm managers in this study were not unusual in the position they occupied. Across professions, new technologies are arriving into expert work that is already in motion – work shaped by long-running pressures from supervisors, peers, clients, and predecessors, and by the small contests over standing that experts navigate day to day. The dominant frame in our field treats the arrival of a new technology as the beginning of the story, with experts' interpretations and responses unfolding from that moment forward. The findings of this paper suggest a different reading. The interpretations experts bring to a new technology, and the practices they build around it are already substantially formed before the technology arrives. What looks like a response to the technology is often a continuation of a performance experts were already engaged in – a way of working through threats they were already managing, using whatever resources the moment makes available.

This reframing matters for how we study technological change in expert work, but it also matters for how we think about the experts themselves. Treating technology as the primary actor and the expert as the responder underestimates what experts are doing when they pick up a new tool. They are not only learning a system or defending a domain; they are continuing the longer project of holding their standing together under conditions they did not choose. The technology becomes part of that project, for better and for worse. Recognizing this is not a softening of the analytic frame but a sharpening of it: it locates the source of variation in expert

response not in the technology, nor in individual disposition, but in the threat landscapes that experts are already navigating when new tools come into reach.

CHAPTER 5: Recognizing as a Recursive Process of Inferring – How Audiences Construct Expertise Inferences Under Digital Visibility

The recognition of expertise is consequential for how people exchange knowledge, ask for advice, and work together. Audiences who correctly attribute expertise to others can seek advice from those most likely to provide it, calibrate their own contributions to what collaborators can do, and coordinate effectively across roles to positive organizational ends (Borgatti & Cross, 2003; Hofmann et al., 2009; Moreland & Myaskovsky, 2000). Audiences whose misattribute expertise risk wasting effort on the wrong sources of knowledge, miss out on useful information, and sometimes damage the working relationships through which expertise is exchanged (Austin, 2003; Su, 2012). Because the stakes of these attributions have measurable consequences for team performance (Lewis, 2004; Ren & Argote, 2011) and social structures in organizations (Leonardi, 2007), scholars have carefully considered the question of how audiences come to recognize expertise in others. The dominant finding in this literature is that experts perform their competence through credentials, demonstrations of skill, and communicative displays that audiences interpret to make their attributions (Abbott, 1988; Collins & Evans, 2007; Eyal, 2013; Treem, 2012). What allows this performative dynamic to work is the structural condition of informational asymmetry. Experts have historically held considerable control over what audiences see. They curate their front stage signals to project competence and hide the backstage activities where the actual messy work (e.g., the deliberations, the false starts, the self-doubt, the mistakes, and revisions) gets done (Goffman, 1959; Pratt, Rockmann, & Kaufmann, 2006; Beane 2019). The opacity of the backstage is not incidental to expert authority. It is a condition under which performers manage other's perceptions of them (Anteby, 2013; Bechky, 2003; Bernstein, 2012; Huising, 2015).

This structural condition, however, is changing. The digitization of work and the platformization of the systems that encode it have generated persistent traces of activity that were once invisible to outsiders, and have made those traces accessible to audiences far beyond the worker (Kellogg et al., 2020; Leonardi & Treem, 2020; Stohl et al., 2016). Patients can now access their medical notes, records, and diagnoses that physicians produce about them. Clients can audit the transactions their financial advisors execute. Even package delivery, automotive repair, and child care services leave digital traces in the form of real-time update emails to track deliveries, digital invoices noting each part and hour of work that went into your car, and app-based records of each time your child eats, sleeps, and gets a diaper change. The unprecedented visibility afforded by these informating, networked technologies perforate the curtain that once separated front stage and backstage work. Consequently, experts who once curated what was visible to audiences are now losing control over how others perceive them.

Yet, it is important to note that visibility is not the same as understanding. Digital traces show what was done as a result of experts' decisions and actions, but not why. They are abstracted from the reasoning, context, and tacit judgment that produced them. Audiences who once saw too little now "see" backstage activity in fine-grained detail without seeing the *meaning* of that activity, and the gap between what is visible and what is interpreted is filled by inference (Leonardi & Treem, 2020). Existing accounts of expertise recognition – built on assumptions of direct or indirect interaction and audience interpretation – do not extend to a setting in which audiences infer qualities about experts from digital traces of work, which carry heavily abstracted signals about expert work. Existing accounts of digital visibility, while attentive to the inferential burden traces impose, have not engaged in detail with how that

burden actually gets met when audiences make inferences about others' expertise. The result is a phenomenon the existing literatures cannot fully explain: audiences are forming durable judgments about expert competence on the basis of partial, ambiguous, decontextualized digital traces, and they are doing so in ways that have substantive consequences for how experts are perceived and how collaborative work transpires.

This paper examines how inferences about expertise change when audiences gain visibility into ambiguous, digitized backstage work. To do so, I draw on a three-year field study of dairy veterinarians in California, who in recent years have gained unprecedented access to new forms of data about cows and farm operations afforded by emerging dairy technologies (i.e., updated herd management software, robotic milking robots, activity tracking collars, and predictive algorithms). Vets and farmers in this industry have historically worked together over decades in close collaborative relationships, with vets gleaning what they could about farmers' practices through periodic visits and direct interactions of farm workers. The introduction of these transformative digital infrastructures into this dynamic gave vets ongoing visibility into a backstage they previously caught only in glimpses, while leaving most of what those traces mean open to interpretation. The findings document a recursive process by which vets constructed expertise inferences from data. From these findings I build a framework of expertise inference under digital visibility that makes three contributions. It theorizes the inferential mechanics through which audiences convert backstage traces into perceptions of expertise. It clarifies a structural shift in the conditions of expert attribution, in which audiences gain interpretive agency while losing the interpretive structure that backstage opacity once supplied, and shows how the relational substrate between observer and observed becomes the apparatus through which interpretation either becomes accurate or solidifies into durable

misreadings. And it extends the visibility literature by foregrounding the inferential work as an analytic object whose components, conditions, and dynamics can be specified and studied. The framework also speaks to the changing terms on which expertise is recognized in contemporary, technology-mediated work, and to what experts can and cannot do about how they are perceived under conditions where the boundary between front stage and backstage becomes porous.

Expert Work Under Digital Visibility

Expertise as Performed and Socially Attributed

Recognizing expertise in others has consequences for organizational function. Research in transactive memory systems (Littlepage et al., 1997), expertise recognition (Libby et al., 1987), and advice networks (Hofmann et al., 2009) all demonstrate that experts' ability to leverage distributed knowledge and coordinate work is more effective when individuals have an accurate understanding of others' knowledge and capabilities. For example, Moreland and Myaskovsky (2000) found that teams who had trained together developed more accurate maps of who knew what and outperformed teams who had received the same individual training. Borgatti and Cross (2003) showed that people seeking advice within organizations chose whom to ask based on their attributions of others' competence, and these attributions shaped both who was consulted and whose advice was followed. Hofmann, Lei, and Grant (2009) found that nurses' decisions about whom to ask for advice depended less on formal credentials than on an accumulated relational sense of which colleagues' judgment could be trusted under specific conditions of uncertainty. The accuracy of expertise attributions matters for coordination, learning, and the flow of advice within organizations, and, as such, the literature has consistently treated this audience-side dimension of expertise as a serious empirical concern

(Anteby et al., 2016; Barton & Bunderson, 2014; Bunderson & Barton, 2010; Collins & Evans, 2007; Treem & Leonardi, 2017).

Given these consequences, scholars have examined how attributions of expertise develop. Expertise is not self-evident. Rather, it is a status conferred by audiences through direct and indirect interaction. Being recognized as an expert depends on the work experts do to signal their competence and on the willingness of audiences to read those signals as credible (Abbott, 1988; Collins & Evans, 2007; Eyal, 2013; Treem, 2012). Experts marshal credentials, deploy specialized vocabulary, perform demonstrations of skill, and manage their conduct in ways that signal that they know things or can do things that others cannot. Typically, these signals are intentionally directed at convincing audiences or passively exhibited as a corollary of performing work. For instance, Treem (2012) showed how consultants in professional service firms communicated their expertise through carefully chosen displays of process and product rather than through any direct exhibition of underlying knowledge. These curated performances (what Goffman calls "frontstage" presentations) project expertise in ways that the performer can manage to present themselves strategically to audiences. Experts have historically had substantial control over what audiences see, and that control has been an essential condition of their authority.

Audiences typically encounter only the polished outputs of expert work while the practices and processes that produced that output are functionally invisible (Nardi & Engeström, 1999; Star & Strauss, 1999). A patient typically only sees evidence of the surgeon's expertise outside the operating room – the appearance of an incision, the surgeon's post-op recommendations, the ease of recovery, and interactions during consultations and follow-up appointments. What they do not see is the private deliberation, the consulted colleagues, the

mistakes made in the operating room, and the host of nurses, technicians, and anesthesiologists who facilitated the actual surgery. Similarly, a client receives the consultant's deck, not the drafts, the rejected analyses, or the scrambling that produced it. A doctoral committee reads the final output of years of study, struggle, self-doubt, and countless discarded drafts in the form of a polished dissertation. This information asymmetry is not incidental to expert authority. Abbott (1988) argued that the inaccessibility of professional reasoning to outsiders is what allows professions to claim jurisdictional control over their domains. Bechky (2003) showed that occupational knowledge is often opaque across boundaries even within the same organization, with engineers, technicians, and assemblers each unable to fully see what the others actually do. The opacity of expert work has given experts substantial latitude to curate their public image. They reveal what supports their claim to competence and withhold what might undermine it which positions audiences at the mercy of what performers are willing to show. The performative work of expertise has therefore depended on the practical impossibility of audiences seeing how that work actually gets done.

Opacity Enables Expert Image Regulation

Goffman (1959) provides convenient vocabulary for this dynamic. He describes the front stage as where social actors perform for an audience, presenting a curated version of themselves consistent with the impression they want to project. The motivation behind such self-presentation is to shape another person's impressions in ways that change their attitudes or behaviors. The backstage, in contrast, is where the work that supports the performance takes place, including the preparation, the rehearsal, the mistakes, and the conduct that would undermine the performance if revealed. These behind-the-scenes activities are often invisible to audiences. Regardless of the modality, every observable action provides information with

which observers make attributions about a particular actor. Performers, therefore, are careful to maintain the boundary between front and backstage so that audiences see only what crosses it. Applied to expert work, the front stage consists of the outputs and visible demonstrations experts choose to show, while the backstage contains the messy work that produces those outputs.

Empirical findings across expert work settings show that opacity is actively maintained by experts as a means of regulating how their competence is perceived. Pratt, Rockmann, and Kaufman (2006) found that medical residents experienced sustained identity work and uncertainty as they grew into their professional roles, but they kept this struggle backstage and presented patients and supervisors with the polished image expected of physicians. The competence patients saw depended on the developmental work happening behind the scenes. Similarly, Beane (2019) showed a related dynamic in surgical training under robotic surgery, where residents engaged in hidden, protocol-violating practice to acquire critical skills that the formal training system did not adequately support. The deviant learning practices had to remain backstage because surfacing them would have invited sanction, even though they were necessary for the residents to become competent surgeons. Outside of medical settings, Huising (2015) showed that compliance officers in research universities built authority by performing visible "scut work" (p. 263) alongside scientists, but the deeper expert judgments they made about regulatory risk remained largely backstage.

Expert authority was constructed through a deliberate mix of selective front – stage visibility and protected backstage opacity. Treem (2012, 2016) found that public relations professionals convinced clients of their expertise through carefully staged displays of process and product even when the underlying substantive expertise was questionable. The polished

front stage performance carried the weight of the attribution because clients had no access to the backstage activity that would have allowed them to assess it directly. Bernstein (2012) made the more general point in his study of a Chinese mobile phone factory, where workers performed less well when their work was fully visible to managers because the opacity they previously enjoyed had allowed productive deviance and informal experimentation that they had to hide under observation. Exposing the backstage did not produce a more accurate view of competence, but it did change what competent work looked like by suppressing the activity that depended on being unseen. Anteby (2013) extended this argument to the institutional level, showing how Harvard Business School maintained a curated image of pedagogy and faculty conduct partly by keeping certain practices off stage.

These studies establish the invisible nature of expert work and how this opacity allows experts to curate their self-presentation. The opacity of expert work enables experts to make claims to expertise that are difficult to scrutinize. It is a condition that experts and expert organizations actively maintain because the polished front stage depends on sloppy backstage activities remaining unseen. The framework is useful here as shorthand for what is visible and what is invisible in expert performances, and the empirical record makes clear that the boundary between the two has been a substantive feature of how experts produce and protect their expert image.

Digital Visibility into the Backstage

Increasingly, this separation of front and backstage is becoming leaky. Two intersecting technological shifts are changing the structural condition that has long supported expert curation. Leonardi and Treem (2020) describe these shifts as characteristics of behavioral visibility – "the increasingly intense amounts of data about people's behaviors that they make

available means that we no longer live in a world where people are invisible and need to work to make themselves visible. Rather, we are always – through traces of us left in data – already visible to others." (p.1605)

The first technological shift that enables visibility is the digitization of work. Activities that once left no trace now generate persistent digital records as a routine byproduct of the tools used to perform them. Leonardi (2014), in a longitudinal field study of a large financial services firm, found that the introduction of an enterprise social media platform made employees' messaging activity – who was talking to whom about what – persistently visible to third-party observers across the organization, transforming previously private communications into a searchable record that others could observe, learn from, and act on. In journalism, newsrooms have adopted analytics dashboards that display real-time reader engagement metrics for every article, making each reporter's audience reach legible to editors and to the reporters themselves (Christin, 2018). In professional services, consulting and law firms have deployed time-tracking, document management, and matter management systems that convert tacit professional judgment into structured data trails that managers, clients, and regulators can audit (Kellogg, Valentine, & Christin, 2020). Importantly, digital traces are generated routinely and incidentally rather than intentionally produced for evaluation. Zuboff (1988) described this dynamic in her account of the informating capacity of digital technology, which produces information about the work as a side effect of doing the work.

The second shift is platformization. Digital records of work are not merely created but made accessible to parties beyond the worker. Shared platforms, dashboards, and integrated information systems give clients, collaborators, supervisors, and external consultants direct access to records that were previously held privately by the worker or their organization. In

software development, platforms like GitHub make programmers' coding activity (e.g., commits, pull requests, code reviews, and abandoned branches) publicly visible to colleagues, employers, and strangers, turning what was once private activity into a permanent, searchable record (Dabbish et al., 2012; Marlow et al., 2013). In many cases, audience members were never intended to see these digital traces and, therefore, "lack the inferential bases for decoding communicative acts of behavior intended for specific targets" (Leonardi & Treem, 2020, p. 1607). The combination of digitization and platformization punctures the boundary between front stage and backstage that experts have historically maintained. Audiences who once saw only the polished outputs of expert work can now see traces of the activity that produced those outputs.

However, visibility into backstage activity is not the same as understanding it because digital signals are abstractions of the work that produced them. They record that something happened without recording why. Leonardi and Treem (2020) develop this point in their treatment of behavioral visibility, arguing that digital traces are simultaneously informative and opaque. The traces show conduct, but the meaning of that conduct requires interpretive work to uncover. A patient can see that their physician ordered a particular test but cannot infer from the order alone whether it reflects routine practice, a specific concern, or a defensive reaction to liability risk. A client can see that their financial advisor sold a position but cannot infer whether the sale reflects a calculated rebalancing, a response to client preferences, or an error in judgment. The signals are available, but they are abstracted from the reasoning, context, and tacit judgment that produced them.

Furthermore, some data quantify outcomes, but not the behaviors and processes that produced the outcomes. This opacity is consequential because digital traces frequently carry

signals about the work that produced them, even when the data were generated for a different purpose. Bernstein (2017) makes a related point about the difference between visibility and transparency, noting that observability of conduct does not entail interpretability of conduct. Christin (2017) shows that the same digital metrics get read very differently by diverse audiences depending on the interpretive frames they bring. Stohl, Stohl, and Leonardi (2016) describe digital visibility as producing information that is available but incomplete, requiring observers to infer what the visible signals imply about what they cannot see directly. In other words, the behavior itself is never visible in data; only sociomaterial performances of it are *implied* in the data, and those performances are abstractions that require inference (Leonardi & Treem, 2020). The inferential work is how audiences construct meaning from signal. Visibility produces inference rather than understanding, and the inferential burden grows with the abstraction of the signal.

The implication is that the structural condition supporting expert image curation has changed without producing the kind of clarity that would make curation unnecessary. Experts no longer have full control over what audiences see. And though audiences can now see much more about experts' work than ever before, they do not have full understanding of the digital traces of work that they now see. Instead, they have visibility into the backstage without transparency into it, and the gap between the two is filled by inference.

The literatures on expertise and on visibility have developed largely in parallel. The expertise literature has examined how experts perform competence for audiences but has assumed conditions of relative backstage invisibility. The visibility literature has examined what digital traces reveal and what interpretive work they require, but it has not engaged directly with how this visibility affects the attribution of expertise. The gap between these

literatures becomes more consequential as the conditions they each assumed begin to shift as digital technologies increasingly permeate expert work while simultaneously providing glimpses into experts' backstage processes (revealing "how the sausage is made"). This paper takes up the question that emerges at their intersection: *How do audiences infer expertise when they gain visibility into ambiguous, digitized traces of backstage work?*

Methods

Research Setting

To examine how audiences infer expertise from digital traces of backstage work, I conducted a qualitative study of large animal veterinarians and their work with dairy farmers in California's central valley. At the time of the study, the California dairy industry comprised of approximately 1,300 dairy farms housing around 1.7 million dairy cows. These operations were legally required to work with licensed veterinarians who oversaw herd health, animal welfare, and public health monitoring. Veterinarians typically visited farms on a weekly or biweekly schedule to perform clinical tasks such as pregnancy checks and ultrasounds, but these visits also served as regular touchpoints for consultation, data review, and problem solving. In this capacity, veterinarians occupied a collaborative role as both service providers performing discrete clinical tasks side by side with farmers as well as consultants assessing the overall health and management trajectory of the herd.

Vets and farmers[1] in this industry typically worked together over years. These sustained collaborations fostered friendly rapport and a deep understanding of each other's preferences, proficiencies, and proclivities as they worked together to optimize herd health, manage

[1] The term farmer here is used to refer to any employee of the farm who performs routine work that is essential to farm operations. Every worker on a farm is required to perform specialized tasks related to their role. I therefore consider them to be experts in their area of work.

operations efficiently, and achieve strategic goals for the business. Farmers relied on vets for clinical expertise, interorganizational insights, and individualized strategic counsel, while vets relied on farmers for insights into the concrete, day-to-day execution of the practices and protocols they developed together. Producing good herd outcomes, therefore, depended on the two parties working in tandem.

Traditionally, vets' views of farmers' work had been largely confined to the front stage. The vet would arrive on the farm, observe what was visible during the visit, examine the cows, talk with the farmer about what had happened since the last visit, consult with them to solve immediate problems, and make decisions and recommendations based on what they could see and what the farmer chose to share. Most of the farmer's actual work – the daily decisions about which cows to treat, which to breed, which to cull, how to allocate labor, how to handle emerging problems – happened backstage, between visits, and out of vets' direct view. Vets caught glimpses of this work through its outcomes, through what the farmer described, and through occasional observations of cows and environmental conditions while on the farm, but the work itself remained largely opaque.

Within this structural separation between front stage and backstage, accurate inference about farmers' expertise were formed primary based on what vets could see and what was performed for them. Vets' effectiveness depended on an accurate understanding of what the farmers knew, what the farmers could execute reliably, and what the farmers prioritized, because every recommendation the vet made had to be calibrated to those features of the farmer's practice. A protocol that worked beautifully on one farm might fail on another because the farmer's capacity to execute it differed, or because the farmer's priorities led them to apply it inconsistently. In short, vets had to make reliable inferences about farmers' capabilities to

apply interventions that would work for their particular farm and work context. At the same time, the working relationship itself was the substrate through which all the vets' work flowed. A vet who upset a farmer through a misjudged recommendation or a poorly calibrated intervention could potentially hurt farm productivity and, ultimately, break the farmer's trust on which the collaboration depended. Accurate inference therefore mattered for two related reasons. It allowed the vet to apply the most appropriate intervention for the problem at hand, and it allowed the vet to apply that intervention in ways that protected rather than damaged their working relationships with farmers.

The adoption of technologies that produced digital traces of farmers' work (e.g., integrated herd management software, activity-tracking collars, and networked information systems) complicated this dynamic. Systems like DairyComp (the dominant herd management system in the dairy industry) digitalized every reproductive event, diagnosis, treatment decision, health outcome, and production metric, producing a durable digital trace of farm activities that could be aggregated, analyzed, and appreciated across time and space. Farmers often granted access to these systems to external experts (i.e., vets, breeders, nutritionists, service providers, business consultants) to improve coordination and collaboration across organizational and professional boundaries, but in doing so, they also provided unprecedented access to backstage farm activity that were previously invisible. Critically, the data they produced were byproducts of farmers running the farms – generated for operational purposes, not as deliberate communications about knowledge, reasoning, or intent – and thus were highly abstract, decontextualized, and ambiguous to the unacquainted viewer.

As such, the dairy context offered a uniquely well-suited empirical setting to investigate how these digital glimpses into the backstage shaped audiences' inferences about performers'

expertise. It is a context in which experts must work together in close collaboration over long periods to examine how inferences shift over time, in which the inferential stakes for the working relationship are high so accurate inferences are essential, and in which digital tools have introduced new visibility into the backstage that audiences must now interpret which sheds light on the mechanisms and nuance of the inference process.

Data Collection

Data was collected between July 2023 and November 2025 primarily through extended interviews with, and participant observations of, twelve large animal vets during their routine farm visits (informants summarized in **Table 5.1**). Farm data and interviews with farmers with whom vets worked closely were also collected from the fourteen dairy organizations (twelve farms and two calf ranches) that vets visited to enhance insights from primary observations and interviews with vets[2]. The researcher's eight-year background working in mixed animal veterinary medicine meant that she could assist in hands-on work (e.g., restraining animals for treatments, drawing blood, assisting in impromptu emergency euthanasia, carrying equipment or labeling sample for processing) and demonstrate general willingness to "get dirty" which facilitated rapid trust-building and legitimacy with informants, bypassing the usual sugar-coating and hand-holding vets described needing to exhibit with inexperienced vet students

[2] As of 2018, California had approximately 7,380 licensed veterinarians (UC Davis, 2020). National data reported by the American Veterinary Medical Association from December 2023 indicated that 4,667 of 82,704 clinical-practice veterinarians were in mixed-animal practice and 3,451 were in food-animal-exclusive practice, together comprising approximately 9.8% of clinical-practice veterinarians (AVMA, 2024). Applying these proportions to California's licensed veterinarian population, and accounting for the dominance of dairy within California's food-animal sector, I estimate the population of veterinarians who practice primarily on dairy cattle in California to be between 100 and 200. By my estimation, the sample of 12 veterinarians represents approximately 6 to 12 percent of this active dairy practitioner population, providing meaningful coverage of a small and specialized professional community.

and externs. This encouraged candid, unfiltered accounts and insights that might not have surfaced with a researcher perceived as a complete novice.

Participant observations took place across naturally occurring work contexts – weekly farm visits, herd health checks, vet-farmer interactions, car rides between farms, lunch breaks, and informal conversations. Each session lasted between four and ten hours, capturing not only formal clinical encounters but also unstructured moments in which veterinarians candidly reflected on their clients, compared farmers, and narrated the reasoning behind their assessments, recommendations, and interventions. During observations, I probed veterinarians' use of farm technologies and data, how these tools affected their interactions with farmers, how they interpreted what data revealed about how the farms operated, and how they identified and addressed various gaps or problems in farm processes. All conversations and observations were recorded with prior consent and transcribed using otter.ai or by hand due to excessive background noise. The researcher also maintained handwritten field notes and took photographs and videos to contextualize audio recordings. At the end of each day, I reconstructed a finely detailed account of the day's activities by expanding handwritten notes into chronological narratives. Then I inserted excerpts from transcripts and indexed photos into the body of the narrative where they appeared in sequence to incorporate as much richness and detail into field notes. Each hour of observation took around ten hours to develop into fully indexed field notes.

Table 5.1. Summary of informant veterinarians.

Practice ID	Informant	Role	Gender	Years in practice	Client farms
VP1	Dr. E	Partner	Female	25	DF6, DF17
VP1	Dr. S	Partner	Female	18	DF6, DF18, DF19
VP1	Dr. B	Associate	Female	3	DF18
VP1	Dr. W	Associate	Female	6	DF20
VP1	Dr. C	Partner	Male	25	–

VP2	Dr. K	Partner	Male	24	DF11
VP3	Dr. M	Partner	Female	21	DF9
VP4	Dr. G	Owner	Male	22	DF12, DF16
VP4	Dr. L	Associate	Female	3	CR1, DF21
VP4	Dr. F	Associate	Female	2	CR2
VP4	Dr. Y	Associate	Male	5	DF16
VP5	Dr. W	Owner	Male	9	DF10

In total, I shadowed eleven vets for twelve days while they performed routine visits to fourteen client farms (dairy farms and calf ranches where dairy calves are cared for until they reach reproductive age). Each morning, I rode to field sites in the vets' work trucks. The drives to and from client farms, lasting between thirty and ninety minutes, were especially fruitful to gain insights into their career histories, their work, and their relationships with their clients, farm workers, and other consultants through unstructured conversations. Once on the farm, I helped prepare equipment and interviewed farm workers who assisted vets during vet checks throughout the day. I observed vets conducting ultrasounds and pregnancy rechecks, routine consultations, and examinations of cows in the pen while recording their conversations with farmers who would scan cows and relay their diagnostic needs from outside the pens. I conducted semi-structured and informal interviews with twelve vets in total. Each interview lasted between one and four hours.

Data Analysis

Motivation for this study grew out of an inductive exploration of technology use in dairy farming more broadly that revealed the critical role of veterinarians and their close interactions with farmers. Analysis followed a grounded theorizing approach (Glaser & Strauss, 1967). Initial open coding of interview transcripts and field notes revealed a pattern around how vets talked about farmers (e.g., "He is a great cow guy…" "He's a rock star…" "He has no cow sense…" "He's a wild card with crazy ideas…"). This prompted me to pay

attention to these explicit evaluations of farmers and how vets arrived at these conclusions about the quality of their knowledge and capabilities. Two patterns emerged consistently across the data. The first was that vets routinely described calibrating their practices to fit with what they believed about farmers' knowledge, capabilities, and value structures across their client portfolios. The second was that vets frequently described moments of surprise in which digital traces of farmers' work contradicted what farmers themselves had communicated or what vets had previously believed about them. These patterns suggested that vets were doing structured interpretive work whose components and dynamics could be specified and whose implications shaped working relationships between vets and farmers.

Equipped with these insights, I turned to the literature on expertise attribution and on digital visibility to sensitize the analysis to existing theoretical frameworks. The visibility literature's account of digital traces as ambiguous and inferential (Leonardi & Treem, 2020) and the expertise literature's account of frontstage performance and backstage opacity (Goffman, 1959; Treem, 2012) helped identify the inferential work I was observing as a form of audience interpretation under conditions where the observed expert had limited control over what audiences saw. This framing sharpened subsequent rounds of coding toward the components of the inferential process (e.g., how vets recognized signals, how they disambiguated meaning of these signals, how they committed to working interpretations, and how they responded to those interpretations in kind).

Next, I developed the analytic categories iteratively through rigorous comparison across cases. Three patterns emerged. First, vets recognized that digital traces of work carried with them implicit information about farmers, but they were not sure exactly what. Second, vets engaged in investigative behaviors to figure out what the signal in the data meant about

farmers' work. Third, vets eventually arrived at inferences about farmers' knowledge, execution, and values that triggered them to intervene. These patterns consolidated into the recursive cycle that organizes the findings section.

Finally, I tested the emerging framework by examining cases that initially appeared to diverge from the pattern. Cases in which vets did not encounter digital traces that gave them glimpses into farmers' backstage activities, in which they jumped to conclusions about what digital traces implied about farmers, or in which disambiguation failed altogether helped clarify the conditions under which the cycle produced productive inferences. For the sake of brevity, I do not explicate these negative cases but instead focus on a clear depiction of the recursive inference process. However, these cases sharpened the account of the recursive structure, particularly the conditions under which inferences emerged and shaped collaborative practice over time.

Findings

Glimpsing the Backstage

Before the widespread adoption of integrated herd management software, activity-tracking collars, and networked information systems, vets assessed farm operations through direct observation and conversations with farmers on the frontstage. They visited regularly, conducted physical exams, asked the farmer questions, and made macro-level judgments based on what they could see during a visit and what the farmer chose to tell them. Inferences about the farmer's competence were made based on direct interactions (e.g., "shit they say" and "how they move around the cows") and physical proxies (e.g., "these cows look skinny today"). The day-to-day decisions about which cows to breed, treat, sell, or cull – and the reasoning behind those decisions – remained largely obscured backstage. Vets caught glimpses of behind-the-

scenes decisions and practices by way of their outcomes but rarely saw 'how the sausage was made.'

The introduction of digital herd management infrastructure changed what vets could see through the dual transformation of digitalization and platformatization. Systems like DairyComp logged every event from the moment the cow was born to the moment they were sold or died (e.g., breedings, diagnoses, medical treatments, injuries, vaccines, calvings). Activity collars tracked individual cow behavior continuously and generated predictions about heats, illnesses, and abnormal feeding and movement behaviors. Networked devices like Bluetooth-enabled ultrasounds and cloud-based databases let vets observe farm activities remotely. Vets could now log in from their phones and review a farm's daily decisions in fine-grained detail, often without the farmers' mediation. Dr. E, a seasoned vet who had practiced for twenty-five years in the region, described the shift as a generational change in what was visible about a farm:

> "The data became more readily available. Twenty-five years ago when I started, there wasn't any of this stuff. There wasn't all this technology. There wasn't this feedback loop that we're getting right now […] There's just more data available. Whereas before, we just didn't know."

This technological transition unfolded incrementally over a relatively short span. Vets described farms moving from physical index cards with handwritten records, to computer-generated lists, to cloud-connected devices that farmers carried in their pockets, and eventually to shared platforms accessible to multiple parties at once – all within the span of thirty years. Dr. E pointed to the most recent transformation as qualitatively distinct from previous waves of technological change:

> "[Farmers] can share [data] with their breeder, their veterinarian, their nutritionist. It's made things more collaborative in that we all have the same information. It's all real time, whereas before, [the data] was more siloed. Now I'm looking at the same graphs that they are. I'm looking at the same trends

> they're looking at – we're all looking at the same trends […] Data has allowed us to critically evaluate the performance of the herd. It forces us to ask more questions, learn the technology, and figure out how to help our clients."

The siloed records of the past had been replaced by shared real-time access to continuous data traces of a farm's operation that no longer relies on the structural brokerage of this information through farmers as previous paradigms of collaboration required. This shared visibility availed forms of evaluation that had not been possible before. Dr. W, a vet of five years, described the depth of inquiry the data now supported:

> "With the amount of data you can get out of DairyComp, there are so many questions that you can answer. Like, how are the cows in pen seven? How are they doing? It seems like they're doing badly, but is it real? I think there's so much stuff that you could learn from the data."

The point vets returned to most directly was about the shift in what was knowable. Before, the questions vets could ask were bound by what they could observe in person and what the farmer chose to tell them. Now, vets could interrogate data independently.

This expanded visibility, however, did not equate to clear understanding about what the signal might indicate about farmers and their work. Vets had to sort signals from noise to extract meaning from what they were seeing. This process operated through three interlinked components – signal recognition, disambiguation, and inference – that cycled continuously, with each component conditioning the others (depicted in **Figure 5.1**). Once vets reached sufficient confidence in an inference about what farmers knew, what farmers could do, and what they valued, they acted on it through calibrated interventions to fill perceived expertise gaps that emerged. How farmers responded to these interventions then fed back into the cycle by generating new evidence that revised subsequent inferences. I present each component in turn, articulate the recursive structure that ties them together, and then turn to response as a downstream practice.

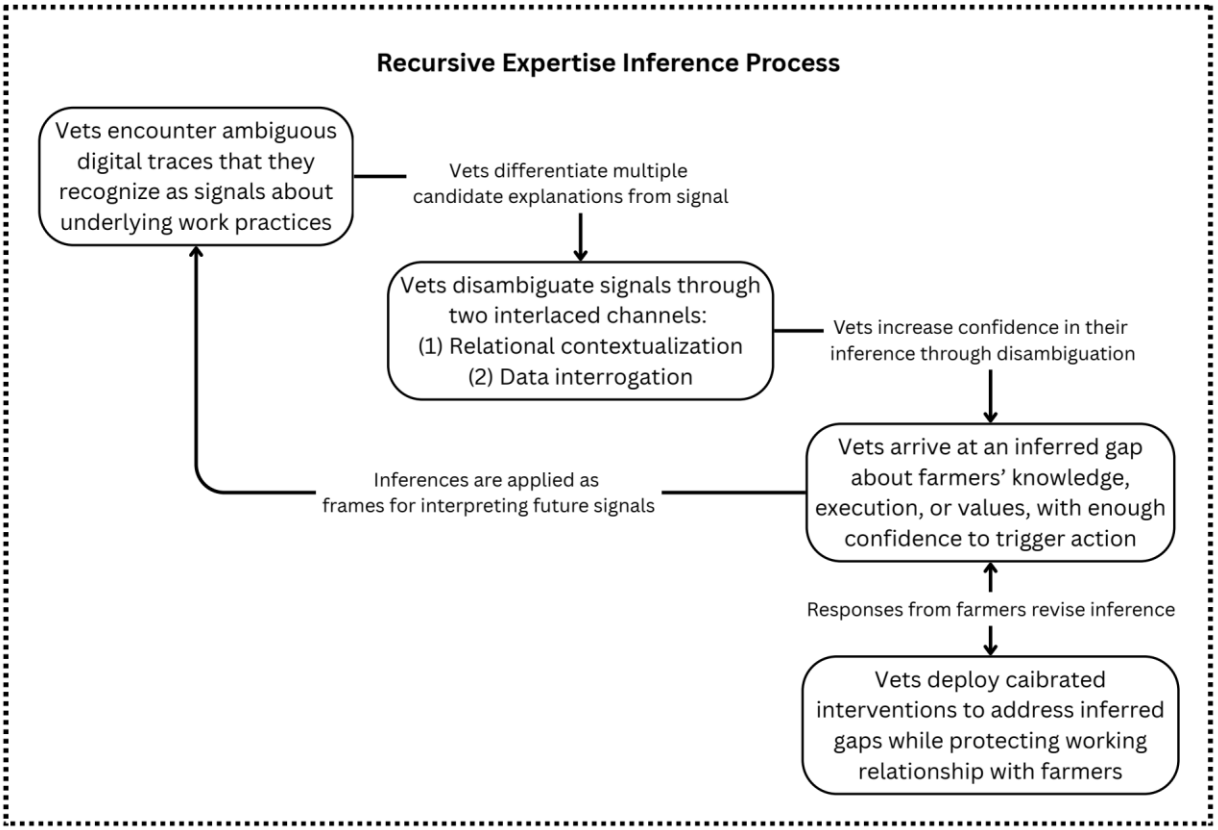


Figure 5.1. A model of recursive expertise inference.

Component 1: Signal Recognition within Existing Interpretive Frames

The digital traces of farmers' work were implied in data about cows (e.g., reproductive activity, treatment records, milk yield, mastitis events, etc.) that were logged as processes and outcomes related to cow care for the operational purpose of running the farm. But vets routinely read these traces on a second register, recognizing data about cows as carrying implicit signals about the people managing the cows. Dr. S, a seasoned vet of eighteen years, recognized this buried signal during a round of pregnancy checks. The data showed an uptick in the number of cows that presented as "open" or not being pregnant when they should have been six months pregnant. The surface-level signal was about the cows potentially miscarrying, but the

unusually high rate sparked her suspicions that something else was at play. What she found when she investigated implicated the farmers' broader practice:

> "They use DairyComp to generate their heat lists. And there was an issue with the collar computer not talking to the DairyComp computer and heats on pregnant cows were not going in. The way I realized that it happened was we were having cows up for confirm check that were supposed to be six months pregnant, and they were open, but [their uteruses] felt like they've been open for a long time. Then I looked those cows up in [the system] and I can see their graphs. I'm like, this cow's had three heats. Why has nobody found her?"

It turned out that the system used to predict heat was not syncing with the system that produced work lists that farmers used to direct their breeding program. Farmers did not notice the syncing issue and, as a result, were not breeding cows that were in heat. What started as a question about why cows were open expanded to a question about why farmers were not catching problems in their digital infrastructure. This expanded into questions about whether they understood how their systems were supposed to communicate and whether the integrity of the breeding program was something they were attending to.

Dr. G, a vet with twenty-two years of experience, recognized a similar opening in an aggregate metric. A drop in the herd's conception rate prompted him to consider what might lie behind it:

> "We're a little off. We have a couple of metrics we check after herd check: percent pregnant. Above 70% is good, we're at 58%. They're not getting pregnant… These cows are a little thin today. It was hot two weeks ago when I was here. I don't know."

The low conception rate signaled that something was wrong but did not specify the cause. What it did tell Dr. G was that something in the farmer's practice or operation needed examination. He pondered the potential explanations for this concerning trend but in the end, he was left with more questions than answers.

Vets recognized that digital traces carried implicit signals about the underlying work that produced it. A signal of this kind opened an array of candidate explanations that vets had to interpret. Many pointed toward the farmers themselves. The data suggested possible gaps in what the farmer understood about the situation, how the farmer conducted day-to-day work behind the scenes, or what the farmer prioritized when making decisions. Dr. E articulated this range of possible interpretations when she described a herdsman whose performance she had begun to question:

> "Luis has been doing this for long enough that he should be able to pick out an animal that is sick before she's dying. Luis has been doing this for 10-15 years. And it just seems like it [lack of ability] could be a number of things. It could be sheer busyness. It could be laziness, or it could just be not having cow sense."

Three candidates sat side by side: being stretched too thin, lacking effort, or lacking the capacity to read animals. Each implied a different root cause for the gaps implied in the data. Dr. F, a junior vet of two years, described a similar range of possibilities at the calf ranch she oversaw, where she found sick calves that had not been treated. The signal was clear (coughing, lethargic calves with no treatment records) but its meaning was not. As she put it, she was unsure "if this means that they [farmers] decided not to treat them, if they missed them, or if they have not walked the hutches yet to do their daily assessments." Each candidate implied a different feature of the farmer's practice – a deliberate decision, an oversight in the daily routine, or a side effect of work design.

Across these cases, signal recognition produced a list of possible explanations but no conclusion. The discipline of holding multiple candidate explanations in mind before jumping to conclusions was itself part of the inferential work. Vets who committed to a plausible explanation too quickly risked misjudging the farmer's practice in ways that could damage the working relationship.

How Certain Signals Become Salient

Signals, while ambiguous, did not occur without prior frames. Vets interpreted candidate explanations against accumulated knowledge of the farmer – what the vet already believed about the farmer's understanding, capability, and priorities. This prior inferential frame shaped which features of the data became salient as signals in the first place. The same digital trace could prompt different recognition in vets who held different prior beliefs about the farmer.

Surprise was the clearest evidence of this conditioning. When a digital trace surprised the vet, it indicated that the prior inferences about farmers' expertise would have predicted something the signal did not match. Dr. E's reflection on one of her well-regarded clients, Steve, captured the dynamic:

> "Sometimes things come up that surprise me, like I think someone is one way, and then you kind of go, 'Oh, I thought you would have known something like that. No? Okay.'"

Dr. E had carried a view of Steve built up over years of working with him. When she encountered a signal that did not fit that inferential frame, the mismatch itself was what made the signal salient. A vet who had carried a different prior belief about Steve – one that anticipated gaps in his understanding – would have read the same signal differently or might not have registered it as a signal at all.

Signals were therefore shaped by prior inferences vets had developed about farmers over the tenure of their relationship. The inferential frames vets brought to interpreting data was itself the product of what they already knew about farmers' knowledge, capabilities, and values. Thus, the inferring process was recursive with each new signal recognized against a frame that earlier inferences had shaped, and each new inference further shaping how future signals would be recognized. But for the signals to produce inferences in the first place, vets

had to narrow their list of possible explanations for the signals to arrive at reliable inferences in the first place.

Component 2: Disambiguation through Two Channels

Once a digital trace had been recognized as carrying ambiguous, implicit signals about the farmer, vets worked to narrow the field of possible explanations through two distinct but interleaved channels. The first I call *relational contextualization* involved situated, interpersonal engagement with the farmer, drawing on relational history vets had built up over years of working together and on direct conversations vets initiated to elicit information from the farmer in the moment. The second I call *data interrogation* which involved independent investigation into farm data (e.g., adjacent metrics, trends over time, comparisons across cows or farms) to look for quantifiable clues as to what is driving the patterns that vets were seeing. Although distinct in mode with the relational pathway operating on the front stage and the digital pathway revealing the backstage, they were entangled in practice. Vets moved between them, used each to inform the other, and developed more holistic pictures of what was happening through their interaction. While vets deliberately engaged in these practices with or without farmers' awareness or intention, these disambiguation channels were contingent upon the quality of the working relationships vets had with developed with farmers.

Whether either channel could be leveraged effectively depended on two conditions that the vet – farmer relationship made possible. The first was the *relationship quality* between vets and farmers which was shaped by positive rapport and trust, and frequency and duration of in – person visits to the farm. Vets worked to build trust with their clients through small, friendly interactions over many years, and frequent on-farm time gave them sustained opportunities to engage one-on-one. Without this engagement, the relational contextualization channel had

little to facilitate it. The second was *access to digital systems*. Without reliable, up-to-date records and the technical permissions to interrogate them, the data interrogation channel was restricted.

These were not practices in themselves but conditions that determined how much information vets could extract from each channel to make sense of signals in the data. When both were present, vets could narrow possible explanations confidently and quickly. When one or both were constrained, disambiguation work became more difficult, and vets fell back on whatever channel remained available – sometimes only on the farmer's own account of what had happened, which was the very account a signal might have called into question. I return to this constraint at the end of the section after presenting how each channel worked when conditions allowed it.

Channel 1: Relational Contextualization

Vets disambiguated signals by engaging the farmer directly, but diplomatically. This engagement involved both accumulated relational history and active in-the-moment questioning. Over years of working with a farmer, vets could typically anticipate their behaviors based on prior responses to advice, characteristic strengths and weaknesses, quirks of personality, family circumstances, and so on. When a new signal came in, vets read it against this profile and engaged the farmer in conversations designed to elicit the additional information needed to disambiguate it. Drawing from their relational histories, vets chose what to ask, what to leave alone, and how to frame their questions to draw out useful disclosures without putting the farmer on the defensive.

Dr. E described a moment when she discovered an unexpected gap in one of her most well-regarded clients. She had long regarded Steve as "top tier" – someone with "tremendous

attention to detail, always looking at data, critically evaluating, and asks really good questions." When Steve began attributing cow losses to bird flu, Dr. E suspected a transition management (supporting cows through calving and post-partum) problem but initially gave him the benefit of the doubt. Rather than challenging the claim directly, she began probing:

> "So now I'm asking innocent questions, like you [a lay person] would be asking innocent questions. I'm like, 'So who diagnoses your fresh pen? Where is your fresh pen?' I'm asking these questions and realizing he doesn't have the answers to these questions... 'Do you think Jose checks for DAs [a torsion of the rumen that commonly occurs during labor], because you have to use a stethoscope?' ... [she mimics his reply] 'I don't know. I don't think so. Do we have a stethoscope?'"

To fill in the context, Dr. E posed her questions with some degree of feigned ignorance. Asking "innocent questions" let her gather information without insinuating that she thought he was lacking in some area. She probed for evidence that would either confirm or revise her suspicion about a transition problem without directly calling him out on his misdiagnosis of bird flu. The conversation disambiguated the immediate signal, confirming that Steve lacked the general awareness of why his cows were dying.

Channel 2: Data Interrogation

Vets also disambiguated signals by investigating the data itself. Where relational contextualization drew on what vets could elicit from the farmer through interpersonal engagement, data interrogation drew on what the digital systems could reveal – sometimes about the cows, sometimes about the work being done with them, sometimes about the people doing that work. This channel often involved investigating the data independently of what the farmer chose to share, looking at associate metrics, expanding the scope of data to look for trends, and comparing what the data showed to what the vet had observed physically. Dr. W, a vet with six years of experience, demonstrated this practice in her remote monitoring of a

farmer's ultrasound technique. The ultrasound was networked via Bluetooth, and Dr. W could connect to his device on her phone without his awareness:

> "Once in a while, I will connect to his ultrasound system on my phone and will watch him scan and catch him missing things. I [monitored] an entire pen, and he didn't put the probe in once!"

Dr. W verified what he was doing in real time and caught him skipping critical steps in his technique that produced unreliable or erroneous pregnancy diagnoses. The digital system rendered his practice visible to her in a way he was unable to curate, and from there she could make sense of discrepancies in the reproductive data.

Many vets used data interrogation to verify diagnostic work by triangulating digital records against physical evidence. Reviewing milk culture results in the herd management system, Dr. E found systematic errors in how farmers were reading plates which had life-or-death consequences downstream:

> "[Holds up a culture plate] This one is not gram negative. She's a no growth. They called 6621 gram negative, and she is a no growth."

The diagnostic codes entered into the cows' digital records did not match what the physical milk culture plate showed. Several cows had been misdiagnosed with intractable bacterial infections and sold as a result. This kind of physical triangulation of data allowed vets to decipher whether the consequence of selling a cow was due to lack of attention, ineffective treatment, misdiagnosis, or judgment error. In this case, it was misdiagnosis.

Data interrogation also operated comparatively across farms. Vets occupied a structural position between farms because they served many clients and saw data from each. Dr. E had formalized this comparative view through a benchmarking spreadsheet she sent to clients monthly:

> "These are all DairyComp [herd management system] numbers so clients can compare themselves to other clients to see where they're at. 'What are you

> doing? What is this guy doing? How come this guy has such a good number on this?'"

Similarly, Dr. W talked about how cross-farm benchmarking forced her to pay attention:

> "I really like it because it makes me look at the data every month too…. I can see if anything jumps out at me, because I know everyone's numbers or about what they should be, or I'll text someone and be like, 'Hey, is this real? Did this dairy really have bad repro?' This farm, for example, when you look at his numbers, he has less than 1% dead, which is pretty good I would say. But if you look at him compared to the average [across farms], he's almost double the rest of our herds… I look at the spreadsheet, and he's double the deads every month. Like, why does he have so many dead cows?"

Dr. W's approach highlighted how the relational and data channels of disambiguation interleaved in practice. She saw cross-farm data and could text her clients to make sense of them. This multifaceted approach allowed her to distinguish individual gaps from systemic ones by leveraging insights from farmers and data. Comparing data across farms pointed to a structural cause rather than to any single farmer's practice, as one vet explained in the case of catching and treating sick calves: "This is an industry-wide problem." When one farm's metrics diverged sharply from others using similar systems, the divergence was itself a signal worth investigating.

Data interrogation had its own internal limits even when access conditions were favorable. Dr. S described a case in which a semen salesman had recommended that a client shorten the voluntary wait period for breeding cows after calving, citing data about improved conception rates related to longer wait periods. Dr. S could not interrogate the data the salesman was drawing on, but she recognized that data of this kind required physical context to interpret:

> "He's only looking at data and not looking at the cows, whereas I am looking at the bigger picture. [The young heifers] are coming in too thin, and then they weren't cycling as fresh cows because they were too thin."

The salesman's data was not wrong, but the inference it supported was confounded by physical evidence the data alone could not reveal. Vets recognized that data had to be read alongside the physical.

Disambiguation Under Constraint

When the enabling conditions were absent or constrained, disambiguation became substantially harder, and vets felt the absence of each channel concretely. Dr. B, a junior vet of three years, described a farm where her ability to interrogate data was restricted:

> "The record keeping has always been terrible. One of the owners just does not reliably use the computer, so readings don't get entered... There's only so much I can do, and they complain about their repro [conception rate], but you don't put stuff in, so I don't even know how bad it is. My hands are tied."

Without reliable data, the data interrogation channel was ineffective. She was forced to depend on what the farmer was willing and able to share, which placed all the disambiguation weight on the relational contextualization channel which was vulnerable to the farmer's own biased and bounded interpretations of the situation.

Relational contextualization with farmers was also constrained if the vet-farmer relationships were not well developed. Junior vets like Dr. L and Dr. F worked with the same farmers as their senior colleagues but with shallower relational archives and weaker standing to probe. They had not yet accumulated the level of understanding and trust that allowed senior vets to read signals against years of experience interacting with their clients. Dr. S, a senior vet, described how this kind of relational context accrued over time:

> "It's just little things over time. Bob [farm owner] is a super nice guy, but Bob doesn't know what the heck is going on. He couldn't log into his TeamViewer. He can't get these reports up. You just figure that out over time through many small interactions. It's the same thing with Miguel [farm manager]. Miguel is an amazing cow man. He can pick out sick cows that I literally walk right by – but he's not a good people person."

This kind of relational acuity was not available to a vet who had been working with a farmer for one year, or who only visited the farm once a month. Without it, signals were harder to disambiguate. Dr. E mentioned this as a factor that constrained her ability to disambiguate Steve's bird flu claim:

> "I only go there once a month. It's my husband's client. We flip flop once a month so he can cover another herd for at least five years now. so, I think I have a pretty good handle on who Steve is, but he's not my every-week client. I just don't know him as well."

The broader pattern was that both relational and digital channels depended on conditions the vet-farmer relationship made possible, and the depth of disambiguation varied with the depth of those conditions. The purpose of disambiguation, when it could be done effectively, was to build vets' confidence in their inferences about the farmer. Working through the two channels narrowed the field of possible explanations to the point where the vet felt confident enough in the most likely cause to commit and act on it. That commitment is the next component of the cycle.

Component 3: Inference as Tentative Commitment

Inference was the point at which vets had done enough disambiguation work to feel confident in a working interpretation of what the signal meant about the farmer's knowledge, capabilities, or priorities. An inference was a commitment to a most likely cause rather than a final verdict. It remained provisional and subject to revision as new information surfaced and as farmers' responses to vets' interventions revealed whether the working interpretation had been accurate.

Confidence in an inference mattered because how vets acted on inferences shaped the working relationship through which all subsequent inference and collaboration would have to flow. Acting on a poorly grounded inference risked harming the working relationship by

challenging farmers' expertise, hurting farmers' egos, wasting time and energy chasing red herrings, producing negative operational or financial outcomes for the farm, or rupturing the trust through which information flowed, negatively impacting the basis on which vets did their work. Acting in ways that harmed this relationship could even result in losing the client altogether. Dr. G put the limit directly: "You gotta watch your ass because you don't want to get fired for pushing too hard on shit." Accurate inference was therefore a precondition for sustainable collaboration. A vet who consistently misjudged a farmer's expertise by prematurely jumping to conclusions about their expertise risked applying interventions that would not fix the underlying issue and would damage the working relationship in the process.

Once vets were confident that they had done enough disambiguation work to form well-reasoned inferences, they fell into one of three kinds of inferences about farmers' expertise: knowledge, execution, or values. Specifically, vets inferred gaps in what farmers knew, what farmers could reliably execute, and what farmers prioritized in decision making. Knowledge inferences implied that the farmer's understanding of the situation did not match what the situation required, and that the farmer did not recognize the mismatch. Execution inferences assumed that the farmer held the relevant knowledge but did not translate it reliably into practice reflecting a gap in doing rather than a gap in knowing. Values inferences indicated that the farmer had the relevant knowledge and the capacity to execute, but the priorities guiding their decisions diverged from what the vet believed the work demanded.

Knowledge Inferences

The clearest manifestations of knowledge gaps appeared when farmers were experiencing sick and dying cows and did not know why. Vets inferred that in many cases, farmers simply did not realize there was a problem until it was brought up. In the case of Steve

and his misdiagnosis of bird flu, Dr. E realized that he was ignorant of the more likely cause, which was a transition problem – "Steve has no idea!" Steve was not failing to follow through on what he knew. He simply did not know about this common issue that was causing his fresh cows to decompensate after calving.

Beyond explicit knowledge, vets frequently cited a gap in tacit or intuitive knowledge about cows as the cause of many oversights. These knowledge inferences pointed to a limit in what vets called "cow sense." Dr. E described it as an innate capacity to "look at a cow and say there's something off about her" that differentiated a "good cow guy" from a bad one. Lack of cow sense was a recurring inference. Dr. E compared two herdsmen who worked at the same farm in similar roles along this dimension:

> "Daniel is the quintessential cow man to the point where he probably doesn't need a collar [heat detection algorithm] to tell him anything, because he's already got it figured out. He can look at an animal and assess in a nanosecond what is going on. And Dylan, he's good at other things, but he's not good at that. I don't think that Dylan will ever reach the level of Daniel [in terms of capability], no matter how many years of experience he has. Some people can look at an animal and see her demeanor – her eyes, her head, how it's held, remember what she normally looks like – and see a change."

For vets, this intrinsic knowing was difficult to articulate, but easy to identify. Dr. L put it more simply: "Good cow guys – they just have it." When the knowledge gap was traced to an absence of cow sense, the inference was that the farmer had reached the limit of what training and experience could achieve.

Inferred knowledge gaps also surfaced as inferences about farmers' meta-cognitive ability to discern good information from bad. Dr. S described this as a common trope among dairymen:

> "There are some dairymen that are very data driven, and they're smart, and they're not gonna fall for bull shit. And then there's some that if you speak to with authority, they sort of just believe you. They call these guys 'the salesman's best friend.'"

The inference was not that the farmer lacked any specific piece of knowledge, but that they did not know what they did not know. They lacked the ability to know when to trust recommendations from various sources and when they needed to scrutinize information before taking them for granted.

A knowledge inference resolved a particular kind of ambiguous signal that pointed to a mismatch between what the farmer understood and what the situation required. Ruling out gaps in execution and values by observing farm practices and probing farmers about their reasoning behind decisions they made clarified that there was a lack of information that prevented farmers from making the correct decisions. Committing to this inference put vets into an informational problem-solving mindset. Their work became about identifying what the farmer did not see and finding ways to surface it. A knowledge inference framed the gap as remediable through learning, which kept the farmer's standing as a competent expert intact and licensed vets to engage more directly to fill the gap with information. Misdiagnosing here was costly. Treating an execution or values gap as a knowledge gap meant pouring time and effort into explanations and reports that would not change the underlying pattern, leaving the actual gap unaddressed and the vet and farmer frustrated.

Execution Inferences

Execution inferences implied that farmers knew what should be done but were not doing it reliably. Dr. E described the general pattern: "I know on any dairy, if you're not checking in, it [process] has drifted off somewhere." According to vets, one of the marks of a "good cow person" was their ability to get stuff done. Dr. E expressed this succinctly after getting off the phone with one of her top clients, Kevin: "He's on it. He's just a man of action."

An inferred execution gap could be observed clearly when Dr. S reviewed the transition cow process at Steve's farm: "They [farmers] are not doing anything correctly… the way they're treating animals, the way they're identifying animals. None of it is being done correctly." Her observations helped diagnose an execution problem with how they were identifying and treating cows that were decompensating after giving birth. However, the underlying causes of procedural drift varied. Some pointed to characteristics of the farmers themselves and their inherent ability to get things done. Dr. E described one client this way:

> "[This client] is pretty severely ADHD, and I think he just really struggles with all the moving parts. With Kevin [a different client], it's just a conversation. It's almost on autopilot sometimes. Over there, I say, 'You guys need to do this right now' and he does it… He's on top of things. At [the other farm], there's always an emergency!"

The inferred execution gap was that this farmer's individual capacity to follow through with many parallel demands was overstretched. Other execution inferences pointed to structural constraints rather than individual capacity. Dr. F described a calf ranch where execution gaps were traceable to how the work was organized:

> "Every couple months, they kind of fall back to their ways of treating animals that are already showing pretty severe signs of disease. Those are the last pens that guys do. They're freaking exhausted by the time they get to those pens."

These workers knew what to do but did not have the capacity to do it by the end of the day.

An execution inference resolved ambiguity in signals by suggesting the relevant knowledge was in place, but the work was not getting done reliably. Ruling out a knowledge gap or values gap through probing farmers to test their knowledge and their underlying motivations meant that vets could see that they understood the situation and had the appropriate priorities but simply failed to execute accordingly. Committing to this inference put vets into a structural problem-solving mindset. Their work became about changing the conditions of the work itself to compensate for missing capacity or perceptual skill. Misdiagnosing here carried

sharper costs. Acting on an execution inference when the gap was knowledge-based wasted the vet's effort on workarounds for a problem that providing farmers with information could have solved. Acting on a knowledge inference when the gap was execution risked signaling to the farmer that the vet thought they did not understand their own operation, which farmers could read as an accusation about competence, and which could narrow what they were willing to share in future encounters.

Values Inferences

Values inferences implied that the farmers had the knowledge and the capacity to execute but that the priorities guiding their decisions did not align with what the vets believed the work demanded. Farmers attended to different things depending on what they believed was important, optimizing for different outcomes and accepting trade-offs that vets perceived as misaligned with the ethos of a good dairyman. Dr. E described understanding farmer motivations as a permanent feature of her work:

> "Every client I have is motivated in a different way, so I don't ever present information or have conversations the same way. Some people just want the facts, or they're interested in what other people are doing, or they're interested in being the best, or they're interested in money, or they're interested in animal welfare – whatever it is."

Inferred values gaps arose when farmers' motivations led to decisions that vets read as harmful to the cows, the herd, or the operation as a whole. The clearest example of this gap appeared when financial priorities overrode other considerations like animal welfare. Dr. F described a calf ranch owner that vets had come to call "the widget guy":

> "He's just looking at the data and the numbers. That's all he cares about. He just wants the data to look good and he wants to make money. He doesn't really care about the calves."

The mortality data on his ranch was visible to him. The data was there. But his value structure weighted financial outcomes above the welfare of the animals. Dr. W described this financial imperative of one of her clients: "[He] needs to break even because the product is somewhat expensive. So he's like, 'I need 19 less dead cows per month to make it break even.'"

Values gaps also appeared as a preference for convenience over rigor, reflecting what vets perceived to be misguided priorities, as Dr. E opined:

> "In the dairy world, dairymen love a quick fix. So, if they can just give a bolus, [mimicking farmers] 'I don't have to do anything else. Yay! Sign me up, I don't care how much it costs.' Even the super smart ones who are current on the latest trends."

The phrase "even the super smart ones" was an inferential distinction. Even farmers who had demonstrable knowledge about cows could fall victim to the alure of shortcuts that minimized hands-on work.

Vets were especially attuned to what these values signaled about farmers' ethical standards. They described a clear delineation between farmers who were "cow minded" (i.e., those who were oriented towards cow health, comfort, and wellbeing as the guiding principle in decision making) and those who were not. Vets treated cow-mindedness as the orienting value of a "good cow person." Dr. S drew on her family's farm to describe what this looked like:

> "My uncle had a dairy, and he was a cow person. His motto was 'the cows come first.' He had it on the wall. That embodies what a cow person is. All their thoughts, all their decision making, all their everything is the cows come first. If you put the cows first, and you always do right by the cows, then the rest of it is going to kind of fall into place."

Dr. M, a vet with twenty-one years of experience, described this guiding principle as one that would enhance farm performance, rather than hinder it: "Just make feed enticing and make life good [for the cows]. I feel like you have to trust a production animal that they will choose the

most productive thing." Comparatively, Dr. B described what the absence of cow-mindedness looked like:

> "They're just so focused on the business mindset that even the people telling them there is something going on with the cows, they're not hearing it."

Dr. G echoed this in more direct terms:

> "You see shit that's screwed up. We gotta talk to [the owner], but he'll be here next month. You're not gonna get shit done… They're not cow people, and they really don't know shit about the cows!"

Dr. G pointed out extreme cases he observed that deeply informed his inferences about farmers' values: "If you see lots of broken tails, lots of downer cows, stuff like that. You start seeing stuff and thinking they don't really give a shit about the cows."

A values inference resolved the residual signal that persisted after information had been supplied and after the conditions of the work had been adjusted. Ruling out a knowledge or execution gap through interrogating farmers' awareness and practices meant that vets could only explain the signals they were seeing by choices the farmer was making about what to attend to and optimize for. Committing to this inference put vets into a relational problem-solving mindset. Their work became about influencing the farmer's priorities without being perceived as antagonistic. Misdiagnosing here was the most relationally risky of the three. Reading a values gap into what was actually a knowledge or execution gap meant questioning a farmer's priorities and ethics over a problem that information or structural support could have addressed, an accusation that could rupture the working relationship outright.

Component 4: Calibrated Intervention

Once vets felt confident in an inference about famers' knowledge, execution, or values gaps, vets faced choices about whether and how to act on it. Farmers were typically unaware they were being silently evaluated so vets had to decide how to intervene to address gaps

without rupturing the working relationship through which all subsequent inferential work and collaboration would have to flow.

The three inference types carried different levels of *relational risk* – the potential to damage the working relationship if not handled tactfully. A knowledge inference carried the lowest risk: it implied that the farmer lacked information, a gap that could be patched without threatening the farmer's identity as a competent operator. An execution inference carried higher risk, implying that the farmer knew what to do but was not doing it – a claim that could be heard as an accusation about reliability or attentiveness. A values inference carried the highest risk, implying that the farmer's priorities and ethos were not aligned with the principles of animal welfare, the most incendiary terrain in animal care occupations. These risks compounded for vets, because the working relationship was the substrate through which all subsequent disambiguation, inference, and collaboration would flow. Misjudging upward (e.g., reading a values gap when the actual gap was knowledge) could rupture a relationship over an issue that a simple education intervention alone could have resolved.

Vets' acute awareness of the compounding risks shaped how they calibrated their interventions through *relational buffering* – a meta-practice that softened the directness of their responses to protect the working relationship from the risk an inference carried. Buffered practices varied with the inference: knowledge inferences could often be buffered through tactful corrections; execution inferences were buffered through structural workarounds that absorbed the gap rather than confronting farmers directly; values inferences required softer strategies still, such as reframing problems, hiding behind data, and in extreme cases, emotional divestment. I treat each inference type in turn, tracing how buffering operated across each, before returning to the escalating dynamic at the end of the section.

Intervening to Fill Knowledge Gaps

The implicit theory underlying a knowledge inference was that the farmer's understanding was incomplete but easily patched. The intervention, therefore, was largely informational. With farmers who took their cues from data, vets surfaced the inference directly. They explained, trained, and produced data and reports to improve awareness. Directness was enabled by the relatively low relational charge of the inference. Not knowing was a fixable condition, and farmers could absorb the correction without their identity or competence being threatened expressly. Dr. E described how she used data as a depersonalized teaching instrument:

> "Data just means we're going to make a decision, not because I said so, but because your data is telling you."

The framing allowed vets to deploy data as an impartial and dispassionate mediator to deliver their recommendations in more palatable packages. Dr. S adapted the format of delivery to the farmer's cognitive style to increase chances of her recommendations landing:

> "If it's somebody that [makes decisions] based on feeling, I'm not gonna email them a spreadsheet. But then I have other dairymen that I absolutely email spreadsheets to."

Dr. E calibrated even further to individual communication styles and preferences:

> "For Derek [owner at DF6], I don't send him a frilly email. 'This is what I think you should do.' Five bullet points. 'Let me know how you want to proceed.' That's it."

For these farmers, the vet identified the gap, named what the data showed, and offered the recommendation in a nonjudgmental and personalized way so the farmer received it in the most effective form.

Not all knowledge inferences could be confronted this directly. When the farmers were invested in their existing understanding (e.g., farmers with decades of experience, farmers

whose self-concept was built around their expertise) direct correction risked triggering defensiveness rather than learning. In these cases, vets shifted to softer Socratic practices designed to let the farmer arrive at the correction themselves.

Data-guided learning. The first practice was *data-guided learning* – letting the data prove itself through repeated correct predictions, allowing the farmer to discover the gap on their own terms. Dr. E acted out this dynamic in her dealings with Paolo, a seasoned herdsman who held strong views about his ability to read cows:

> (Acting out a typical exchange) "[I would say,] 'She looks good. Do you want to wait? Or do you want to treat her?' He would opt to wait and I would come back the next day – 'Is she sick?' (mimicking Paolo:) 'Oh, yeah, she has pneumonia.' You almost have to let the data show them, 'Oh, maybe I don't know,' because me saying, 'You're gonna follow this list, and I don't care what you say,' is not going to work."

Paolo arrived at the inference about the limitations of his own knowledge himself, on his own terms, without Dr. E confronting him with his ignorance directly.

Strategic acquiescence. A second practice was *strategic acquiescence* – agreeing with the farmer's proposed course of action even when the vet disagreed privately and waiting for the evidence to materialize. Dr. E described this logic in how she approached an especially self-determined client:

> "There's a certain amount of strange power in just saying, 'Yeah, you want to try that? Yes, let's try that.' And in my head I'm thinking 'worst idea ever,' but then I think of ways to maybe share some advice. And then you wait, and you wait for your opportunity to point it out and say, 'Oh, look how this uptick [in mastitis rates] just happened.'"

The farmer was allowed to experiment. The vet held the inference privately, watched the evidence accumulate, and surfaced the correction under the cover of observation only once the data had made the case. Dr. W described the same practice as relational investment:

> "I think he [Alex, farm manager at DF20] just needed someone to go along with his crazy ideas. After we kind of did a few things where he was crazy, and I just kind of let it go, I think he just started trusting me."

Acquiescence built trust that could later be used to confront more consequential issues.

Jedi mind tricks. A third practice was what Dr. L called the "*Jedi mind trick*" – planting an idea so that the farmer would arrive at it themselves and adopt it as their own:

> "You say something along the lines of, 'I wish there was a way that we could do this.' And they follow your logic, and they're like, 'Well, what if we do this?' – 'Oh, interesting. Yeah, that will probably work. Maybe it's worth a try.' And then they start rolling with that."

Dr. M summed up this strategy: "It has to be their idea." The farmer experienced the correction as self-generated rather than externally imposed. The vet's role in producing it was hidden, deferred, or laundered through strategic influence. Dr. E described how she would later credit the farmer when the data confirmed her own private suspicions:

> "I made sure that the data supported their idea. I was like, 'Oh, you were totally right. Look at that. Now we know that this is the best way to go with [the data].' And I try to reference it back. It's like positive reinforcement."

Dr. B described a more passive version of the same practice:

> "Convincing him [owner at DF18] to do anything is not an easy task, and I usually have to just lightly float the idea a couple times and then wait for him to bring it up on his own."

Farmers typically responded positively to these interventions since vets made it appear to be their own ideas or conclusions even though they were surreptitiously orchestrating them.

Intervening to Fill Execution Gaps

When the inference was about execution, the intervention could not be informational, because lack of information was not the gap. Vets intervened functionally instead by changing the conditions of the work – taking on more supervision themselves, absorbing functions the farmer could not perform reliably, or letting technology perform the task – redistributing the

burden of doing across more entities. Each of these practices avoided the relational risk of telling the farmer that they were unreliable or incompetent.

Intensified monitoring. The first practice, *intensified monitoring*, involved the vet's continued attention to whether the work was getting done. Dr. E described how this had become a routine part of her role on certain farms:

> "They're failing at execution and it's my job to make sure the big stuff gets done. I can't fix everything, but I can make sure that there's follow through on the big stuff. I just check up on it."

Dr. L described a similar role at the calf ranch (CR1) she advised:

> "My job is mostly to monitor the calf doctors [a colloquial title for calf caretakers] to make sure the herdsmen are following their directives."

Monitoring took a structural form rather than a confrontational one, thereby buffering potential risk by folding this practice into vets' farm visit routines imperceptibly.

Ad hoc role expansion. The second practice, *ad hoc role expansion*, involved the vet absorbing some high-level functions the farmers could not perform themselves. When a farm lacked someone with the technical skills to manage and use the digital systems, the vet would step up. Dr. S described her own role as a *de facto* technology manager at some of her farms:

> "I do the computer work for two of my dairies because they don't have anybody that's good at it, and it happens to be something that I am good at and I enjoy, so I kind of step up and do it."

Dr. E captured the same dynamic more bluntly: "Some days our job is just to show up and call DairyComp for them." These were not explicit roles or services that farmers expected from their vets, but where a gap could be filled by the vets, they absorbed these tasks to help their clients' farms run more efficiently.

Technological substitution. The third practice, *technological substitution*, involved deploying technology to fill the gaps in execution. Vets described facilitating implementations

of digital tools like health detection ear tags to compensate for practical constraints. Dr. E described this in the case of Dylan, a herdsman whom she had inferred struggled with detecting sick calves, who could rely on illness detection ear tags for calves:

> "If you have somebody like Dylan who struggles with [finding sick calves], in a position like that, technology fills in the gap… The list is there. These are the animals I need to look at. They probably look at that list like it's gospel. Because she's on the list, even though her eyes are bright, I'm gonna maybe just give her some vitamin B12. She's in the system now… That's where the collars can help them, and whether that's because you're gone from the dairy and you don't have time to walk those animals, or your guy just can't learn it properly, either way."

Technology filled execution gaps so that farmers could focus their attention and effort in ways that meant they could do more, intervene sooner, and achieve better outcomes for cows.

Intervening to Fill Values Gaps

When the inference was about values, the gap was not addressable through information nor through structural workarounds alone. The gap reflected what the farmer chose to address and what they prioritized. Vets intervened by finding ways to influence farmers' decision-making priorities without confronting the values themselves. Direct confrontation about values was the most risky intervention that could rupture the working relationship so vets incorporating strong buffering strategies.

Reframing issues. The first practice was *reframing issues* – translating problems and recommendations into the language of what the farmer cared about. The widget guy at Dr. F's calf ranch (CR2) did not respond to appeals about animal welfare. He responded when financial outcomes were threatened: "He gets upset when he starts not making money." Dr. E described how she had come to organize her practice around this kind of translation work:

> "It's no longer just a health thing, now it's a business thing. And you have to think about the money they have – do they want to grow [the herd]?… All of the things."

Dr. B made the same explicit reframing:

> "You have to gear your mindset for each dairy you step onto. If these guys don't care about that, we don't need to talk about that."

Reframing acknowledged the values gap internally without confronting it externally. Vets understood what farmers valued and tailored how they presented the problems to align with existing value structures. The vet did not say "you should care about cow welfare more." They said "this practice is costing you money," and the recommended action – the same action that would have served cow welfare – was adopted on financial grounds. Vets leveraged data to strengthen these reframings, as Dr. E explained:

> "Now, the veterinarians in the world that are working with these farmers are breathing a sigh of relief, because we can just say, 'Let's look at the data. What does the data show?' They're taught that data shows you things, whereas before they trusted their own perceptions, which are biased, and never really looked at the facts… Data is powerful, because [I can say], 'Fine, don't believe me. Believe the numbers.'"

The data became a way for farmers to appreciate the consequences of their decisions in a framework that resonated with their values.

Strategic detachment. Some farmers, however, refused to be swayed. When problems could not be reframed in ways that produced changes in farmers' behaviors, vets resorted to *strategic detachment* to preserve their working relationships. Rather than continuously banging their heads against a wall, vets eventually realized that their sustained efforts yielded little to no positive result. After vets exhausted their options in their journey to fill inferred gaps in expertise, they eventually concluded that farmers values were simply not aligned with what they believed should be the guiding ethos of someone in their position – and could not be reconciled. Dr. S described this decision to detach when reflecting on interventions that led nowhere:

> "I advise. I use the data to help me advise when I can... and then the cards are out there. That's my opinion. And then you guys do what you're gonna do… I think if you get too involved in other people's problems, and they become your own, it's just a lot of stress."

The farmer had the information, had the capacity, and was choosing not to act on it. By the end, the inference had escalated to a gap in values, and Dr. S's mindset shifted accordingly: "You gotta pick your battles." All the vets shared this experience. They continued to advise but invested less in whether the advice was followed. Dr. G described his threshold directly:

> "You gotta know when to lay off. I'm gonna harp on you two or three times about something. Then after that, if you don't want to do it, I don't care."

Strategic detachment was the limit case of intervention and relational buffering. The inference remained unaddressed. The gap was not closed. But the working relationship was preserved through the vet's emotional withdrawal – but only to a point. Vets withdrew to preserve their own mental wellbeing and energy, but in doing so, their perceptions of farmers often suffered in tandem as vets' inferences about them turned to personal failings.

The Recursive Inference Cycle

Signal recognition, disambiguation, and inference operated as a continuous cycle. New signals were recognized against the interpretive frame established by prior inferences; disambiguation was directed by possible explanations gleaned from signals; and inferences were revised as disambiguation evidence accumulated. The cycle had no terminal endpoint at which the vet was "done" inferring about a given farmer. Each new visit, data review, and conversation generated material that fed back in – including, critically, how farmers responded to vets' efforts to fill the gaps vets had inferred.

This recursive structure had two consequences for how vets came to know the farmers they served. The first was that inferences were continuously evolving, with each response to a

gap-filling effort returning new evidence to the cycle. Vets typically defaulted to the knowledge inference as a starting point unless disambiguation pointed clearly toward an execution or values gap. It was the most generous interpretation, the easiest to address, and the least costly if wrong. Vets revised upward toward execution and then toward values only when more generous interpretations were ruled out, typically through the failure of an informational response to change behavior. Dr. E's escalating inferences about Steve illustrate this dynamic across encounters:

> "The first month that he said something to me, I told him something's wrong… I took a milk sample of his tank and sent it off for bird flu. It was negative... This is not right. We sent Dr. S out there and she went through the fresh pen treatment protocol with the guys. And she was like, they're not doing anything correctly… he definitely heard us and he addressed the things that he liked, but he didn't fully address the whole thing, because he's still hanging on to the bird flu thing… He doesn't even really see the urgency… I went back the next month. And I was like, 'How are things?' And he's like, 'I still think I have bird flu!'"

Dr. E started by ruling out a knowledge gap, testing for bird flu, and receiving a negative result. When that failed to resolve the problem, she escalated to an execution gap, sending another vet to verify that the fresh cow protocols were not being followed. When Steve still clung to his bird flu theory after this intervention, Dr. E revised again, this time toward values, attributing his behavior to a failure to see the urgency. At each step, the farmer's response to her gap-filling strategy fed back into her inference about what he knew, how work was being done, and what he prioritized. What began as a knowledge inference revised toward execution and eventually toward values as repeated correction failed to change behavior.

The second consequence was that inferences could become sticky. Once a vet had committed to a view of a farmer's knowledge, capabilities, or priorities, that view tended to persist and shape future signal recognition. Dr. E's reflection on Steve captures this dynamic:

> "I wouldn't have known that about Steve. I would have assumed that he's just a rock star. He is a rock star in most areas. This is his weakness, and now I know that so in the future, when he says something random, I know that now I gotta dig. I gotta ask a few more questions. I'm not gonna take him at face value in that situation, because I know that he's not very good at piecing out where the problem really is, and that's my job now."

The accumulated stock of prior inferences functioned as an interpretive lens through which subsequent signals were filtered. This stickiness made inferences efficient to apply, but it also meant that vets carried interpretations forward that subsequent evidence might not have warranted on its own. A farmer who had been read as lacking cow sense in one cycle would have new signals read against that frame, even when those signals might have supported a different reading on their own.

The cycle therefore did two things at once. It produced working inferences about farmers' expertise that vets used to inform their interventions, and it accumulated those inferences into the frames that would shape what counted as signal in the future. Each failed response did more than revise the working inference – it escalated that inference from the relatively neutral terrain of knowledge gaps toward the more relationally fraught terrains of execution and, eventually, values.

Discussion

The findings of this paper showed how external experts construct interpretations of another expert's competence from digital traces of work the other expert did not produce for evaluation. Drawing on the empirical case of dairy veterinarians who had gained access to digital records of farm operations (i.e., medical treatments, reproduction events, activity data, and production metrics), I documented a recursive process by which vets recognized digital signals carrying implicit information about the farmers' actions, disambiguated those signals through situated engagement and direct data interrogation, arrived at working inferences about

the farmer's expertise gaps along three dimensions – knowledge, execution, and values – and acted on those inferences through calibrated interventions whose directness varied with the relational risk the inference carried. Across all three inference types, vets operated through a meta-practice of relational buffering, strategically softening their interventions to protect the working relationships through which all subsequent collaboration and inference was possible. The cycle was recursive in two ways. Prior inferences conditioned what counted as signal in subsequent encounters, and farmers' responses to interventions fed back into the cycle to revise inferences over time. See **Figure 5.2** for theoretical model.

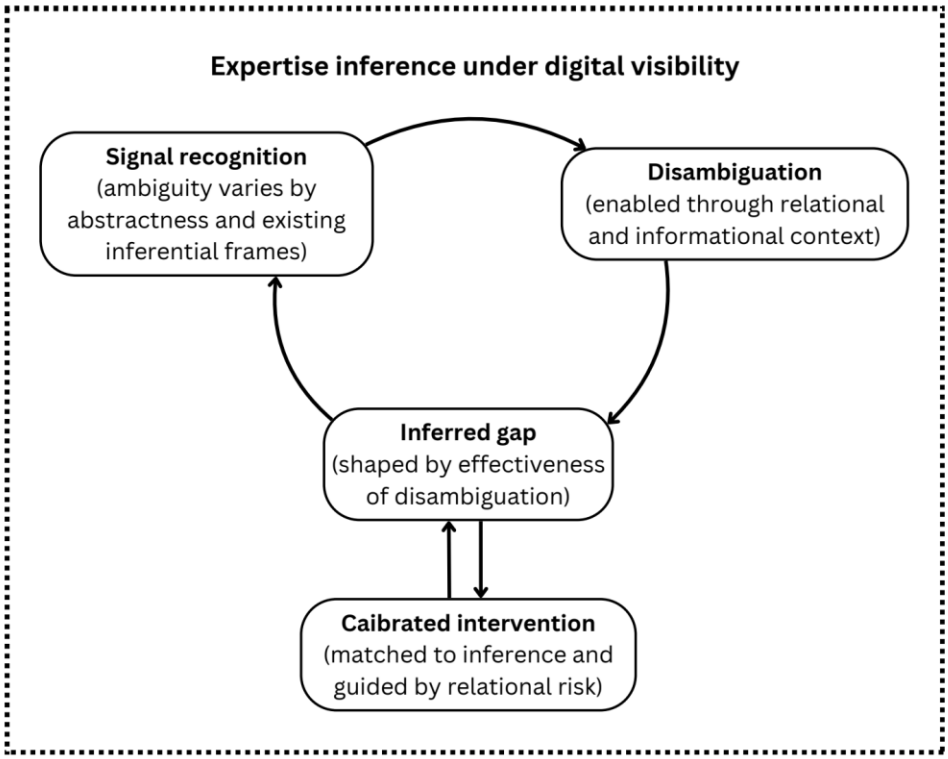


Figure 5.2. Model of expertise inference under digital visibility.

This pattern speaks to three theoretical conversations. First, it advances a framework for understanding expertise inference under conditions of digital visibility, generalizing the empirical findings to contexts where digital traces of decisions, actions, and outcomes

permeate collaborative expert work. Second, it has consequences for the expertise literature, particularly for how scholars theorize the conditions under which experts can manage their own attribution and how audiences form attributions about experts. Third, it extends the visibility literature by foregrounding the inferential burden visibility places on observers and the dynamics through which that burden is met.

A Theory of Inference Under Digital Visibility

The expertise literature has long understood that audiences attribute competence to experts based on what experts make visible to them. Goffman's (1959) front stage and backstage distinction has provided convenient vocabulary for this dynamic. Experts perform on the front stage through credentials, communicative displays, and demonstrations of skill, while the backstage of actual work – the deliberation, the false starts, the routine decisions – remains opaque to outsiders. Audiences attribute expertise based on what crosses the boundary between the two, and the boundary itself has historically been controlled by the expert. The expertise literature has examined the front-stage side of this dynamic in considerable detail (Treem, 2012, 2016; Bunderson, 2003; Collins & Evans, 2007), documenting how experts perform competence and how audiences interpret those performances. The visibility literature has documented the structural shift that digital tools and platform-mediated work have introduced into this dynamic, showing that the backstage is increasingly leaving traces that audiences can see but that are increasingly abstract and ambiguous as a result (Leonardi & Treem, 2020; Stohl, Stohl, & Leonardi, 2016). What this literature has not specified is what audiences do with the backstage once they can see it. Visibility is not the same as understanding, and the inferential work of converting traces into inferences about expertise has remained largely unexamined.

The contribution of this paper is a framework that addresses this gap. The framework theorizes audiences as actively constructing expertise inferences from backstage traces through a recursive interpretive process. This recasts what the front stage and backstage distinction does under digital visibility. The boundary between the two no longer separates what audiences can see from what they cannot. It separates what experts can curate from what audiences must infer. This is a fundamental shift in the function of that boundary. In the classical account, the front stage and backstage operate as a visibility boundary. Experts manage attribution by controlling what crosses it, and audiences form judgments from what experts choose to show. The boundary acts as an opaque curtain, and the expert decides what passes through. Under digital visibility, however, the boundary still exists, but it is no longer opaque. It functions as a veil that distinguishes between two kinds of material audiences encounter while not fully offering clarity: the curated displays the expert has shaped for interpretation, and digital traces the expert generated but never intended audiences to see. Under digital visibility, both are visible to the audience, but only the first carries the interpretive scaffolding the expert built into it. The second arrives stripped of that scaffolding (no context, no reasoning, no insights into what produced it) and must be made sense of by the audience on its own. The expert's control over attribution narrows accordingly while the interpretive burden on the audience broadens. Experts retain authorship over signals they curate but lose authorship over signals they unwittingly generate, and the share of what audiences see in the backstage grows as digital infrastructure deepens. This backstage activity becomes partially visible through the veil and yet remains fully opaque about the reasoning and practices behind those traces, and audiences must work across this gap to construct meaning. The interpretive work the framework documents is the work of constructing that meaning, and the recursive structure is how the

meaning gets built up, refined, and sometimes solidified into enduring frames about the character of one's expertise that govern subsequent interpretations.

Two features of inference under digital visibility distinguish it from the inferential work the expertise literature has examined. The first is that signals are *byproducts* of work. The data traces observers see are generated as the expert does their work, but do not convey any information about their practices or choices directly. The expertise literature's accounts of attribution have focused on settings where experts produce signals for audiences and audiences interpret those signals as information about what they know, what they do, and who they are (Treem, 2012; Bunderson, 2003). When audiences encounter backstage traces, they are reading information the expert generated for operational purposes – to log a treatment, record a transaction, complete a task – without curation or context. This changes the inferential task. Audiences must reconstruct the reasoning behind these ambiguous, decontextualized, and incomplete trace without the framing that there is information to be gleaned about expertise at all. In other words, data traces can be so abstract that the process through which observers produce meaning from them requires a considerable amount of interpretive and inferential work in and of itself. This is the condition Leonardi and Treem (2020) describe when they argue that digital traces require inference in proportion to their abstraction from the work that produced them. The framework developed here specifies the inferential mechanics through which inference under high levels of abstraction actually gets done.

The second feature is that digital visibility produces conditions under which audiences gain interpretive agency but lose interpretive structure. The recursive cycle of signal recognition, disambiguation, and inference can, in principle, run with little to no input from the observed expert if operated through data interrogation alone. When audiences have

substantial access to digital traces of an expert's work but limited opportunities for direct engagement with the expert, they construct expertise inferences primarily through their own interpretation of abstract and ambiguous backstage traces with limited resources to disambiguate. Agency shifts from expert to audience. Audiences decide which traces count as signals, which to disambiguate further and how, which to read as evidence of competence, and which to read as evidence of gaps. This is interpretive agency the audience did not previously have. Under direct interaction, the expert organized what audiences saw through curated performances that supplied the interpretive structure within which audiences' attributions took shape (Treem, 2012, 2016). The performance itself framed the meaning of the signals it conveyed. Under digital visibility, audiences encounter signals stripped of that framing. The signals are byproducts of work that the expert produced for operational purposes, not communications designed to guide interpretation. Audiences therefore have wider latitude to interpret, but they also have less to anchor their interpretations against. This is a structural shift in the conditions of attribution. The expertise literature has assumed that expert self-presentation organizes audience interpretation (Treem & Leonardi, 2017), but this framework shows that under digital visibility, much of the interpretive structure that performance once supplied is no longer constraining and audiences can construct their own.

These structural consequences compound over time. Inferences formed early in the expert – audience relationship become the interpretive frame through which subsequent signals are read. As the cycle iterates, the frames reify. Audiences carry forward views of the expert that become increasingly difficult to refute or escape unless evidence emerges to destabilize them. This is consistent with broader accounts of how prior beliefs shape attention (Weick, 1995), but it has particular force under digital visibility because the interpretive structure expert

performance once supplied is no longer the preeminent influence. Furthermore, inferences can become enduring over time. Sticky frames are the temporal residue that adhere the observer to versions of the expert that were inferred during prior cycles of inferential work and steer downstream disambiguation strategies and interventions. The longer audiences observe an expert, the more substantial the frame becomes. And because the expert was not privy to its construction, they may have neither awareness of the frame nor access to the conditions under which it could be revised. Bernstein (2012) showed that workers responded to increased visibility by hiding the practices that opacity had previously protected. The framework here shows that these defensive moves may not work under conditions of backstage visibility but the consequences of inference still very much affect the experts being evaluated.

The framework also clarifies that effective inference under digital visibility depends on a solid relational substrate. The agency audiences hold over interpretation does not, on its own, produce accurate inference. The two disambiguation channels through which audiences narrow candidate interpretations both depend on relational conditions that expert and audience produce together. Direct engagement requires the relational access to ask questions and the standing to probe. Independent interrogation of data depends on the working relationships through which audiences are granted access to the systems where traces live. These channels supply the contextual information audiences need to constrain their interpretations. Without them, audiences still form interpretations, but the interpretations are made from a thinner base of evidence and are more vulnerable to misdiagnoses or the recursive solidification described above (Shami et al., 2009). The relational substrate, in other words, restores some of the interpretive structure that digital visibility has undermined. It is what allows audiences to disambiguate ambiguous signals against the contextual knowledge that direct engagement and

trusted access provide. Huising (2015) showed that experts build relational authority through the visible work they do alongside the parties who evaluate them. The framework here extends this logic. The relational conditions experts and audiences jointly maintain serve as the apparatus through which audiences' interpretive agency can produce accurate inferences that preserve the integrity of the relationships themselves.

The Changing Conditions of Expert Attribution

The expertise literature has long understood that expertise is performed and attributed rather than simply possessed (Goffman, 1959; Treem, 2012; Collins & Evans, 2007). Scholars have documented the substantial work experts do to manage their self-presentation by curating what audiences see (Anteby, 2013; Pratt, Rockmann, & Kaufmann, 2006; Treem, 2016), and the literature has converged on the view that expert image management depends on a structural condition: audiences see only what experts choose to make visible. The findings of this paper show that this condition is changing in ways that the expertise literature has not yet engaged.

When digital traces of expert work become available to audiences who did not encounter them through the expert's curation, the conditions under which expertise is attributed shift in three substantial ways. The first is that experts lose much of their control over how they are interpreted. The image management strategies the literature has documented assume a structural separation between the front stage of curated performance and the backstage of actual work. Experts manage attribution by choosing what crosses that boundary. The findings of this paper show what happens when the boundary itself becomes porous. Vets attributed expertise to farmers not through the displays farmers chose to perform but through traces of routine work farmers produced for operational rather than performative purposes. Farmers had little ability to anticipate which features of their work would be interpreted as expertise signals

or to curate those features as they would have framed a deliberate performance. The findings of this paper show that when traces experts leave become visible, the consequences extend beyond "who knows what" to the broader question of who controls expert attribution.

The second shift is that experts are increasingly evaluated based on partial and ambiguous digital traces rather than based on situated performances designed to convey competence. The literature on expertise recognition has shown that the accuracy of attribution depends on the richness of information available to audiences (Bunderson, 2003; Borgatti & Cross, 2003; Hofmann, Lei, & Grant, 2009). Direct interaction provides audiences with opportunities to ask questions, probe, and update their judgments in real time as the expert clarifies, demonstrates, or corrects. Inference from digital traces operates under abstract conditions. The traces show what was done but not why, and observers must construct an account of the reasoning to attribute expertise at all. The interpretations audiences form can be confident even when the conditions for confident interpretation are absent, and the recursive structure of inference means that early interpretations shape later ones in ways that can entrench rather than correct misreadings.

The third shift is a consequence of the first two. The dimensions of expertise that show up in digital traces become more salient in evaluative contexts than dimensions that do not. Tacit judgment, embodied skill, and relational competence – dimensions central to expert practice (Collins, 2010; Hadjimicheal et al., 2024; Polanyi, 2009) – are difficult to capture in digital traces, while dimensions that are easily logged become readily available for evaluation. When attributions are increasingly made from traces, the dimensions of expertise that are captured begin to define what counts as expertise in digitally-mediated evaluative contexts.

The expertise literature has yet to engage with how expertise is recognized and attributed when evaluative conditions become digitalized.

These shifts together describe a phenomenon the existing literature on expertise cannot fully explain. The classical conditions of expert attribution (e.g., curated performances and communicative exchanges) are not the conditions under which much expert evaluation now operates. Audiences see, and infer from, digital traces of work that experts did not produce for them and may not even know are being read. Experts can no longer rely on the front-stage management strategies the literature has documented because the front stage no longer fully mediates audience access. Understanding how expertise is attributed under these conditions, and what experts can and cannot do about it, is a substantive theoretical task the literature has not yet taken up. The framework developed in this paper documents how the loss of curatorial control and the inferential burden placed on audiences have consequences for how expertise is inferred and attributed, and for how experts work together under these new evaluative conditions.

Inference as Recursive Work

The visibility literature has long understood that digital traces require interpretive work to be meaningful. Scholars have established that digital traces are abstracted from the activity that produced them, often opaque to outsiders, and rarely transparent in the sense of conveying what they mean without inferential effort (Bernstein, 2017; Leonardi & Treem, 2020; Stohl et al., 2016). They have shown that audiences read digital traces against interpretive frames they bring to the encounter, with the same signals producing very different conclusions across observers (Christin, 2017). They have demonstrated that audiences form inferences that are sometimes accurate and sometimes not, and that inferences treated as accurate by audiences

shape behavior regardless of whether they actually are (Marti & Gond, 2018). The literature has, in other words, established that inference is a constitutive feature of digital visibility and has begun to characterize its properties, its conditions, and its consequences.

What remains underdeveloped is a structured account of how the inferential process unfolds. The literature has identified that audiences infer (Leonardi, 2015), that their inferences depend on the abstraction (Stohl et al., 2016), and that the resulting interpretations have downstream effects on expert work (Hofmann et al., 2009). What audiences do in the interpretive moment – what they recognize as a signal, how they disambiguate it, how they commit to a working interpretation, how they revise that interpretation across encounters – has been treated as an important feature of visibility without being examined in depth. The paper's first contribution to the visibility literature is to develop the inferential process at this level of detail. The recursive cycle of recognition, disambiguation, inference, and intervention specifies what audiences do step by step, identifies the structural conditions on which each step depends, and shows how the steps feed back into each other across time.

This explicit process makes possible a more granular engagement with questions the literature has raised but not fully resolved. The accuracy question, for instance, has typically been examined as an outcome – does the audience's inference match the underlying reality? – rather than as a function of the interpretive process that produced it (Barton & Bunderson, 2014; Bunderson & Barton, 2010). The recursive cycle reframes accuracy as a property of how the cycle was conducted wherein inferential accuracy is the cumulative product of a structured process whose components can fail or succeed at different points. This reframing has consequences for empirical work. Studies that ask whether inferences are accurate can begin to ask where in the cycle accuracy was produced or lost.

The paper's second contribution to the visibility literature concerns the conditions under which audiences can do this inferential work effectively. The literature has examined how interpretive frames vary across audiences (Christin, 2017) and how cultural and institutional contexts shape what gets seen and how (Flyverbom, 2016). This study highlights the relational and structural conditions between observer and observed that enable or constrain effective disambiguation. My findings show that the depth of inference depends on conditions that are jointly produced by audiences and the experts they observe – direct engagement that allows audiences to probe and clarify and access to data systems that allows audiences to interrogate digital traces against context. These are properties of the working relationship within which inference takes place. Without them, audiences still construct interpretations, but the interpretations are made from a thinner base of evidence and are more vulnerable to error and to recursive solidification. The depth of audience inference depends on the same relational conditions that scholars of professional work have shown matter for advice-seeking, knowledge-sharing, and authority-building (Hofmann, Lei, & Grant, 2009; Huising, 2015). Visibility scholars examining inference can productively integrate these conditions into their accounts.

A third contribution, briefly noted here and developed further in future work, is that the framework opens up cross-domain expertise inference as a phenomenon that the visibility literature has not yet fully engaged. Most existing work on inference under visibility has examined within-domain cases – colleagues reading colleagues, members of the same profession or organizational work groups interpreting each other's cues. The dairy case examined here is structurally different. Vets are external experts whose own professional frame overlaps with but does not fully match the farmer's, and the digital traces they read were

produced within the farmers' unique knowledge schemas, workflows, and logics. Reconciling this variation involved accurately inferring the meaning of signals so they could apply interventions designed to protect and preserve their working relationships. Cross-domain inference of this kind is increasingly common as platforms make work across professional boundaries visible to evaluating audiences in adjacent domains, and it raises questions about how the recursive cycle operates when the observer and observed do not fully share the interpretive frames.

Conclusion

The dairy veterinarians and farmers in this study are not unusual in the position they occupy. Across professions, digital tools are giving collaborators ongoing access to one another's work in ways that previous generations of practice did not allow. What once happened backstage is now logged, transmitted, and read across distances and roles. This shift is often described as a question of efficiency, oversight, or coordination. The findings of this paper suggest it is also something more fundamental. As technology gives us new visibility into one another, it changes the basis on which we come to know, evaluate, and trust the people we work with. Judgments about expertise – and the working relationships those judgments shape – are increasingly built from traces rather than from direct encounter, and the inferences we form from those traces accumulate into images of one another that the people we are imagining have little ability to participate in producing. The veterinarians in this study did this inferential work with care, in long-standing relationships they had reason to protect. Many contemporary settings offer fewer of those protections. As the conditions under which experts are perceived continue to shift, understanding how that perception is constructed – and what it

costs when it goes wrong – matters for the quality of collaborative work and for the integrity of the relationships through which that work is done.

CHAPTER 6: CONCLUSION

This dissertation began with a theoretical puzzle. The scholarship on expertise has been fragmented across three ontological camps that have largely avoided one another. A knowledge-based camp, rooted in cognitive science and psychology, treats expertise as something experts possess and asks how new technologies change what experts know and how they come to know it. A performance-based or situated camp, rooted in the sociology of expertise, treats expertise as something experts enact through communicative and embodied displays for audiences and asks how technologies reshape what experts do. A recognition-based camp, rooted in the sociology of professions, treats expertise as something audiences confer through processes of legitimation and asks how technologies reshape the criteria by which expertise is evaluated. Each camp has produced an impressive body of empirical foundations. But each also treats its preferred conceptualization of expertise as if it could be analytically isolated, with technological change cast as a disruption that produces effects readable within its own conceptual vocabulary.

The three studies in this dissertation were designed to probe these dimensions empirically through a single case study of the implementation of data-intensive technologies into California dairy farms. Study 1 examined what happened to knowing when farmers came to engage their cows through continuous digital representation. Study 2 examined what happened to performing when farmers facing pre-existing threats to their expertise interpreted a new technology through the lens of those threats. Study 3 examined what happened to recognition when audiences (i.e., veterinarians) gained digital visibility into the backstage of farmer work.

What these findings reveal collectively is that knowing, performing, and recognizing are not three separable phenomena. Rather, they are three facets of one process. When technology enters an expert work system, it reconfigures all three at once, and the effects are amplified because the three processes are entangled with changes in one feeding back into the others. Over time, the feedback loops produce trajectories of expertise that no single mechanism, considered in isolation, would predict. The aim of this conclusion is to highlight the entangled, relational dynamics through which technology co-constitutes expertise by drawing up an integrated framework that can inform studies of technology, expertise, and organizing.

Before laying out the integrative model, it is worth reflecting on what each chapter teaches us that is difficult to synthesize when read individually. Each chapter, while focused on a particular constitutive component of expertise, contains within it traces of the others – moments at which the analysis brushes up against the other two processes, like tectonic plates colliding and transforming the topography and trajectory of both. These collision points are fusion sites for new and fertile analytical ground.

Study 1: Knowing as Embedded in Performance and Recognition

Study 1 examined the representing process through which farmers came to know cows under quantification through digital technologies. Its central finding was that representing was an ongoing process in which the same physical cow was continuously rendered as several different kinds of objects (a body that did or did not fit the robot at the inscription bottleneck, a biographical record of past events, an algorithmic prediction of future states, and a position within a comparative distribution) each of which invoked its own evaluative logic and its own definition of what it meant to be a "good cow." Farmers worked across the friction between

these renderings and their physical animals through deliberate effort to align digital and physical realities (what I call *congruence work*). Where the representation was more plastic and editable, farmers modified the digital. Where the representation was rigid (e.g., a robot arm with fixed architecture, a proprietary algorithm, a herd-level average), farmers modified the cow.

Read on its own, this study is about technology changing what experts know. But the knowing it documents cannot be separated from performing and recognizing. As farmers came to know cows through digital representations, what they could perform to audiences shifted with it. Demonstrating sensory skill at the cow's side was no longer sufficient. Farmers were now expected to be fluent in net merit indices, sick chance, mastitis likelihood scores, and probabilities, to know cows in the new register of formalized metrics and to make that knowing visible to veterinarians, breeders, peers, and consultants in novel and unfamiliar ways. As farmers navigated performing new forms of knowledge, audiences began to recognize them for knowing new kinds of things about cows – not the cow in the pen, but the cow as she appeared across the digital renderings the technology produced.

The physical object of a cow herself was also changed through this process. Each rendering invoked new evaluative criteria – for udder conformation, for genomic merit, for behavioral compliance – and farmers translated those criteria into culling and breeding decisions that, over years, reshaped their herds. The object of knowing was transformed both epistemically, as it became legible in new ways, and physically, as artificial selection narrowed morphology, dispositions, and genetic baselines. What farmers knew about cows had to evolve alongside the physical evolution of the cows themselves.

Recognition then fed back into knowing. As farmers were recognized for competence in algorithmic interpretation, data-driven breeding, and predictive intervention, this recognition created pressure to keep developing knowledge in those directions rather than others. Knowing how to interpret information that sensors could codify became the baseline as veterinarians gained access to the same data. Visibility came as a doubled-edged sword providing farmers with deeper insights into farm operations while giving the same access to those who evaluated their use of those insights. These findings therefore show more than that technology changed what farmers knew about cows. It shows that knowing reorganizes what knowledge experts can draw upon to perform expertise for audiences, how experts are able to demonstrate knowing, as well as how these performances will be perceived.

Study 2: Performing as Translating Knowing to Recognition

Study 2 examined how four farm managers interpreted and took up the same technology in radically different ways. The central finding was that the technology entered ongoing expertise performances that were already under varying degrees of threat from non-technological sources, and farmers' interpretations of the technology were filtered through what I call a *lens of threat* defined by the intensity and specificity of the threats they were already managing. Farmers' orientations toward new robotic milking systems tracked the threat conditions they were already navigating at the moment it arrived. The farmer whose expertise was uncontested perceived the technology itself as a threat and adopted a critical orientation marked by scrutiny and competition. Those facing pre-existing threats from other sources – doubts from superiors, dismissal by outsiders, insubordination from workers – interpreted the technology as a resource for managing those threats and adopted deliberative, deferential, or

appropriative orientations accordingly, each using different affordances of the same system to address the particular vulnerabilities they faced.

Read on its own, this study is about technology changing how experts perform expertise. But findings suggest direct linkages to knowing and recognizing. The threats that constituted the lens through which farmers interpreted the technology were consequences of recognition dynamics being read by farmers as evidence of how audiences had come to doubt their expertise. The orientations farmers adopted in response then reshaped what they came to know. Some were driven to develop data interrogation skills to scrutinize algorithmic predictions while others were prompted to lean into their cow sense to differentiate themselves. Other farmers chose to divert effort towards limiting access to data or appropriating the symbolic value of the technology to bolster their status and authority. Each orientation directed attention toward different features of the work and produced different bodies of knowledge and skill as a result.

What farmers came to know from these different orientations then became the foundation of expertise upon which audiences recognized them. Traditional dairy expertise had been visible in the body. A farmer walking the barn, palpating an udder, or running a hand down a cow's flank produced an embodied display that audiences could see and read as competence in real time. Interpreting an algorithmic flag, weighing a conception rate against a herd-level average, or deciding whether to override a threshold happened in the farmer's head. The work was cognitive, often silent, and produced no outward sign that anyone could attribute to expertise without explicit framing. The new performance paradigm meant that farmers had to find new ways to make expertise legible to audiences. Technology forced them to develop new ways of performing what they knew in ways that audiences could appreciate.

These findings are, in effect, an extended demonstration that performing is the hinge at which past recognition dynamics shape what experts orient toward, what they come to know, and how they convert that knowing into displays audiences can perceive. Recognition pressures from the past condition the orientations farmers adopt toward the technology. Those orientations shape what farmers come to know. And what they come to know creates new front-stage conditions and new props that experts must learn to incorporate into their performances in ways audiences will read as expertise.

Study 3: Recognition as Both Reading and Conditioning of Expert Work

Study 3 examined how dairy veterinarians constructed inferences about farmers' expertise from digital traces of farm operations that farmers had not produced for evaluative purposes. The central finding was that recognition under digital visibility operates through a recursive interpretive cycle wherein vets recognized digital signals as carrying implicit information about the farmer, disambiguated those signals through situated engagement with the farmer and independent data interrogation, committed to working inferences about the farmer's knowledge, execution, and values, and acted on those inferences through interventions calibrated to relational risk. Prior inferences conditioned what counted as signal in subsequent encounters, and farmers' responses fed back into the cycle to revise inferences over time. The chapter showed that the frontstage/backstage distinction that has organized scholarship on expert impression management no longer separates what audiences can see from what they cannot. It separates what experts can curate from what audiences must infer.

Read on its own, this study is about technology changing how audiences recognize and interpret expertise signals. But the recognition process depends on knowing and performing. The digital signals vets read (e.g., reproductive events logged in DairyComp, drops in

conception rate, milk quality reports, the gap between when a cow had a heat detected and when she was actually bred) were traces of farmers' work that were shaped by what they knew and how they acted upon it. A farmer who deferred to the technology as an act of symbolic appropriation was producing a different pattern of decision traces than a farmer who was scrutinizing the technology to extend and demonstrate his interpretive skills. Audiences reading these traces were reading the cumulative residue of how knowing had been reorganized and how performing had been adopted under technological change.

Inferences about expertise also shaped how vets perceived subsequent performances. Actions of a farmer the vet had come to view as a strong interpreter of data would be interpreted differently from those of a farmer the vet had come to view as deferring too readily to the technology, even when both were doing similar work in front of the vet. The interpretive frame audiences carried about a farmer's expertise determined what counted as a genuine demonstration of competence and what counted as misjudgment or as use of the technology as a crutch. The reading depended less on the act than on the accumulated frame through which it was being interpreted. This had consequences for what farmers could effectively perform. Performance moves that fit the frame were registered as expertise whereas performance moves that ran against the frame were either not registered at all or were read as confirming a deficit the vet had already inferred. Recognition frames therefore did not just sit downstream of performance; they organized in advance what performance could accomplish.

Audiences' inferences about expertise also fed back into knowing through calibrated interventions. Vets who believed farmers lacked certain knowledge or skills actively sought to fill those gaps, shaping what farmers knew and what they could do. A farmer recognized for data fluency received less support in this domain than a farmer whom vets viewed as lacking.

A farmer recognized as a "cow guy" rather than a "data guy" and was positively reinforced to continue enacting the embodied skills that recognition rewarded would invest more heavily into retaining that image. Over time, what farmers knew tracked what their audiences had been recognizing them for, which itself had been shaped by what they had been performing.

These findings show that recognition is not a downstream consequence of expert performance but a process that runs in both directions. Expert knowing and performing are shaped in advance by anticipations of how audiences will judge them. And the frames audiences hold about experts organize how that knowing and performing will subsequently be read.

Integrating the Three Studies

Rather than thinking of expertise as a thing that technology either preserves or displaces, this dissertation conceptualizes both expertise and technology as continuously shifting, malleable processes wherein technology and expertise are interdependent, entangled, and mutually constitutive. As Pakarinen and Huising explain, technologies are "embedded in the network of interactions through which the relational expertise of professions is constituted" (2023, p. 1). Orlikowski and Scott further argue that "an important challenge for research going forward is developing ways of thinking and talking about the social and material worlds as inseparable, as constitutively entangled" (2008, p. 463) – a challenge this dissertation takes up by specifying how technology becomes endogenous to expertise across three interrelated recursive processes.

Existing research has shown that knowing, performing, and recognizing each influence one another, but the literature has not developed a unified account of how they interconnect, and the linkages between these domains remain undertheorized. Barley's (1986, 1990)

influential work on CT scanners in radiology departments comes closest to explicitly theorizing these linkages, showing how radiologists' initial lack of operational knowledge fundamentally altered their epistemic relationship to medical imaging, which in turn transformed how they performed their expertise in relation to technicians, ultimately reconfiguring the professional hierarchy and status relationships between the two groups. Barley's analysis demonstrates that technology does not change one dimension of work in isolation; rather, technological change sets in motion transformations across multiple interconnected dimensions of expert work.

Building on this insight and drawing from sociomaterial perspectives that conceptualize technology as part of rather than separate from social processes, this dissertation proposes an integrated framework of technology-expertise co-constitution. Rather than treating technology as an exogenous force that disrupts, automates, augments, or reconfigures an existing form of expertise, this framework positions technology as an endogenous participant in three interrelated processes through which expertise is constituted.

Read together, the three studies show that what looks like expertise in any moment is the joint product of three processes – knowing, performing, and recognizing – that condition one another continuously. Each study established that its focal process is itself recursive: representing reshapes the object of expertise, performing responds to and reshapes ongoing threats, and inferring builds and revises judgments about competence. Technology participates in each of these recursions as an endogenous element.

The three recursions feed into one another. Study 1's representational work produces the data traces that vets read in Study 3. Study 2's performance orientations shape which traces get produced and which get withheld. And the inferences vets form in Study 3 feed back into

farmers' future work in two ways: through the interventions vets make, which alter the threat conditions that shape farmers' next performance orientations, and through the evaluative frames vets come to carry, which orient what farmers come to know next. The system is one composite loop with three interleaved processes, and technology has entered each of them at once. See a model of this depicted in **Figure 6.1**.

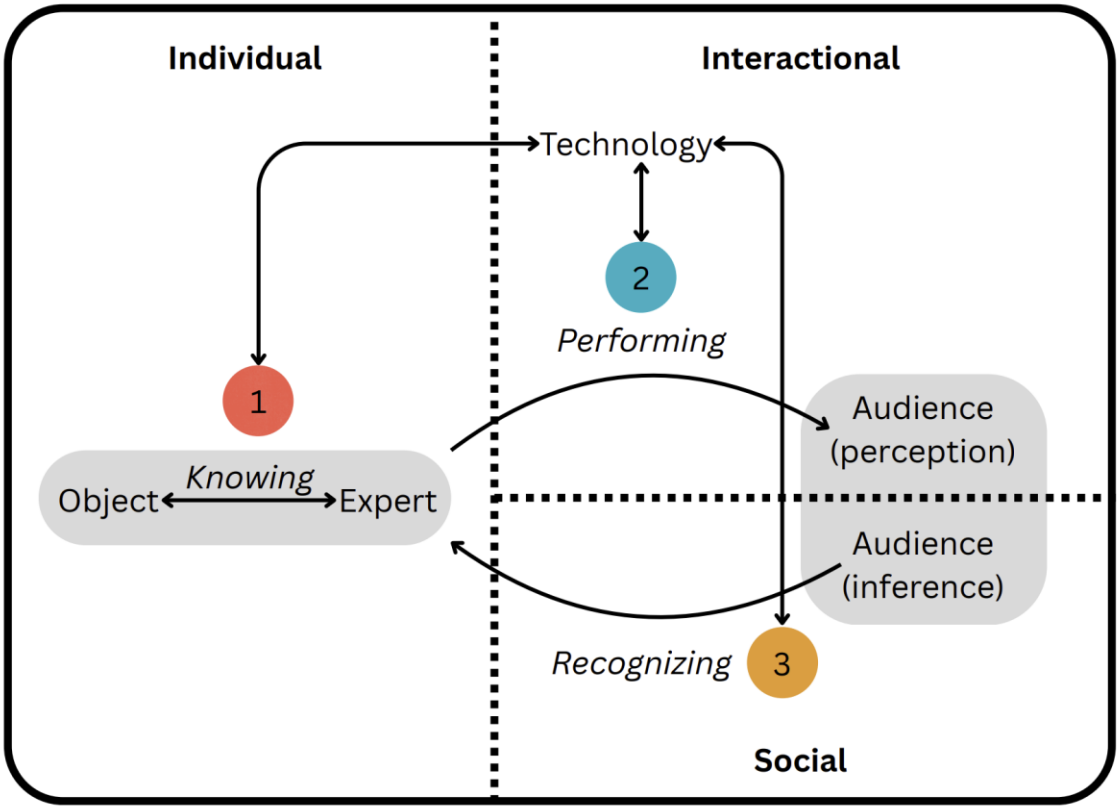


Figure 6.1. Integrated model of the three interrelated recursive processes that constitute expertise, and how technologies become endogenous to each.

The Limitation of Single-Process Accounts

The dominant approach in studies of technology and expertise has been to treat one of these processes as the analytic starting point and the others as effects. This narrow scope is in part necessary to produce tractable studies, but it understates the complexities of how

technological change moves through expert work. When the three processes are interconnected, examining only one means observing a small part of a larger system and attributing to that part what is being produced by the system as a whole.

This limitation manifests across several traditions. Studies of technology-induced automation and augmentation (e.g., Brynjolfsson & Mitchell, 2017; Frey & Osborne, 2017; Raisch & Krakowski, 2021) treat technology as a substitute or complement for expert *knowing* and ask what experts will still be needed to do. The framing brackets the performance and recognition dynamics through which "what experts do" is socially established in the first place, leaving the analysis unable to account for why technically substitutable work is often not actually substituted. Sociomaterial studies of technological change (e.g., Orlikowski & Scott, 2008; Leonardi & Barley, 2010; Faraj & Pachidi, 2021) have insisted on the inseparability of the social and the material but have made uneven progress in operationalizing this insight through specific empirical phenomena; expertise, treated as one of those phenomena, offers a way to specify what the entanglement actually consists of. And studies of technology and expertise more broadly (e.g., Bailey et al., 2012; Beane, 2019; Lebovitz et al., 2022) typically focus on one of the three processes and treat the others as background. The integrative framework does not invalidate these studies. It situates them, and it reveals what the single-process framing leaves out.

Reinterpreting past studies of technological change in expert work illustrates what an integrative reading can surface. Barley's (1986) CT scanners are conventionally read as epistemic disruption that produced structural change; an integrative reading shows the shifts in knowing, performing, and recognizing happening together, conditioning one another in real time. Nelson and Irwin's (2014) librarians are read as a sequential identity transformation; an

integrative reading shows the identity moves had purchase because they were paralleled by simultaneous reorganizations of what librarians knew and how they performed value to students and faculty. Anthony's (2021) financial analysts are read as a case of performance preservation through black – boxing; an integrative reading shows that black-boxing also reshaped what analysts came to know about the algorithm and produced output patterns that audiences later read as evidence about competence. Bechky's (2020) forensic examiners are read as a recognition shift driven by DNA evidence; an integrative reading shows the new standards reshaped what examiners had to know and how they had to perform to remain credible. Across these cases, the integrative framework does not displace the original analyses; it situates them, surfacing the dynamics each was not equipped to theorize from within a single-process frame.

Expertise as a Relational Ecology

The model treats expertise as a relationally constituted ecology in which three processes – knowing, performing, and recognizing – operate together and condition one another. Knowing is the expert-object relationship (how experts interpret, classify, and intervene upon their domain), performing is the expert-audience relationship (how experts render competence visible to others), and recognizing is the audience-expert relationship (how others evaluate those displays and confer or withhold expert status). What I call "expertise" is the ongoing achievement that emerges as these three relations co-evolve. These processes are nested, in that each contains conditions originating in the other two; entangled, in that movement in one means movement in the others; and recursive, in that each cycle produces outputs that become inputs to the next, with recursions crossing between processes rather than staying internal to any one. The system is not three independent loops but one composite loop

with three interleaved phases, held together by what I call relational contingencies – the jointly produced configurations of expert, object, audience, and technology that determine which paths the recursive cycle can take. Technology participates in each relation as an endogenous component.

The implication for studying technology and expert work is direct. When technology enters this system, it does not enter at a single point. It enters all three relations at once, and because the relations condition one another, its effects propagate rather than remaining contained within the dimension a researcher happens to study. Its impact is more than the sum of its functional, relational, and symbolic effects considered separately. Robotic milking systems were, on their face, a functional change in how cows were milked, but the inscription bottleneck narrowed herds and changed what farmers came to know about cows, reorganized farmers' daily routines and how they explained their work to others, and changed what veterinarians read in the data and inferred about farmers' competence. There is no clean way to say the technology started in one domain and produced consequences in the others. The changes happened together. What looks from any single vantage point like a contained effect (e.g., a productivity gain, a workflow change, a status shift) is the visible surface of a reshaping that moves through all three relations at once.

Studying expertise as an ecological system has consequences for how we approach technology and expert work. It means treating the expert not merely as a container of knowledge and dispenser of skill whose expertise is being acted on by technology, but as a participant in a relational system whose position is being reshaped across multiple dimensions at once. It means treating the worker not as a node to be augmented or automated but as a

person embedded in working relationships, evaluative communities, and material environments whose configurations together determine how their expertise can manifest.

This is partly an analytic commitment and partly a humanistic one. Studies that examine technology's effect on one dimension at a time (e.g., what it automates, how it shifts performance, how it changes relations) capture real phenomena but only from one angle. They miss the complex ways these effects compound and condition one another in the worker's experience. A change to what an expert is known for is also a change to what they are seen doing and to what they will come to know next. The worker living inside this system experiences the technology's arrival as a reshaping of their whole orientation within their network of relations, not as an isolated change in any one part. Acknowledging the relational and ecological character of expertise grants both the complexity of expert work systems as relational infrastructures and the humanity of the worker whose experience is shaped by all three processes at once.

The contribution of the integrative model, then, is not only theoretical. It is an argument for viewing expert work holistically as a complex, embedded, multidimensional achievement – and for studying technology's effects on that work with the same complexity.

Implications

For the Study of Expertise

For expertise scholarship, an ecological ontology offers a way past the fragmentation that the three camps have produced. The dominant move in this literature has been to take expertise as an analytic object and ask which of the three ontologies – substantive, performative, or social-constructivist – best captures it. The integrative model suggests that narrowly focusing on one excludes important dynamics of the others. Expertise is not one kind

of thing whose nature can be settled by choosing the right ontology. It is the ongoing achievement of a relationally contingent loop in which three processes operate together, and the dominant ontology in any moment of analysis depends on which phase of the loop one happens to be looking at. The expertise scholar who studies knowing is not looking at a different thing than the scholar who studies performing or recognizing; they are looking at the same phenomenon from different angles. The model does not eliminate the distinctions between these camps. Rather, it explains why each captures something real but partial, and what each misses when it works alone. The implication for future research is that studies of expertise can fruitfully be designed to follow technology's interaction across multiple processes where they might uncover hidden mechanisms at play. However, the methodological costs of this approach are substantial – it requires multi-perspectival, longitudinal designs that track experts, their objects, their audiences, and their evaluators together – but the analytic payoff is access to the entanglement and amplification dynamics that the single-phase studies cannot capture.

For the Study of Technology and Organizing

For studies of technology and organizing, an ecological ontology contributes to the long-running effort to operationalize sociomateriality through specific empirical phenomena. Calls for endogenous theorizing of technology in organizational processes have proliferated, but the empirical traction has been uneven. The model offered here operationalizes sociomateriality through the specific lens of expertise construction: it specifies how technology becomes endogenous to organizational processes not by adding it as a variable to existing models but by showing how it participates in each phase of a recursive loop through which experts and audiences continuously reconstitute one another. The implication is that the

question scholars have been asking – how does technology change expert work? – is a limited starting question. In contemporary organizations where technology is ubiquitous in expert work, a better question is how technology participates in the continuous constitution of expert work, and what happens to the ecological system when technology becomes a participant in multiple processes simultaneously. This reframing has consequences for how empirical studies of technological change in expert domains are designed. While most studies of technology and work compare differences between pre- and post-implementation settings to attribute the changes to technology, the integrative model suggests this design will systematically underestimate the entanglement and amplification dynamics that produce the observable outcomes, because it sees the outcomes without seeing the recursive loops that produced them.

For Practice

For practice, an ecological ontology has consequences for how organizations think about introducing new technologies into expert work. The standard implementation playbook treats adoption as a problem of skill development, behavior change, or culture change, and addresses whichever domain its proponents are trained in. The integrative model suggests that all three need attention simultaneously, and that interventions in one domain that ignore the others will produce trajectories the organization did not intend. Fiona's case is again illustrative: the organizational logic that put the technology in her hands and let her decide how to use it produced an appropriative orientation that, over time, narrowed her practical expertise and weakened her standing with the very audiences she most wanted to convince. An implementation approach that had attended to the performance conditions she was navigating – and that had supported relational contingencies between her and her workers, her veterinarian, and her external consultants – might have made different orientations more

approachable. The practical implication is not that there is a "right" way to implement technology in expert domains. Simply reskilling or upskilling workers will not address the ways audiences perceive them if they are unable to perform expertise in ways that were formerly recognized. Implementation choices that attend to only one of the three processes – knowing, performing, or recognizing, in isolation without attending to the broader ecological system – will be amplified into trajectories that are difficult to reverse. In sum, organizations must adopt technologies prudently and consider its effects on the ecological balance of their workers.

Limitations and Future Directions

Several limitations should be acknowledged. The dissertation works within a single empirical context – California dairy farming – and while this context is well-suited to the theoretical task because it makes all three processes visible at once and creates conditions under which they can be observed in relation, the model's portability to other expert domains is difficult to confirm empirically. Domains where the object of expertise is less physically present (financial assets, software code, medical images), where audiences are less proximate (regulatory bodies, distributed clients, anonymous platform users), or where digital visibility takes different forms (audit trails rather than continuous sensor data) will likely produce variations in how the loop operates. Comparative work across domains will be necessary to specify the boundary conditions of the model.

A second limitation concerns the temporal aperture of the empirical work. Although the data span three years and include multiple cycles of observation, the trajectories the model theorizes operate over decades. The orientations farmers adopted to the technology in 2022 will produce trajectories whose full shape will not be visible until 2032 or beyond, and the

changes in cow physiology and herd composition that congruence work is producing operate on similar timescales. Longitudinal extensions of this work, ideally returning to the same farms over many years, will be necessary to specify how the recursive cycle unfolds across longer arcs.

A third limitation is methodological. The model places considerable weight on relational contingencies – the joint products of experts and others that govern how the loop operates – but the methods used in this dissertation, while well-suited to observing these contingencies, are not designed to vary them systematically. Future work could productively combine the ethnographic approach taken here with comparative designs that vary the relational substrate (e.g., farms with strong veterinarian relationships versus weak ones, organizations with rich peer networks versus thin ones) to test how trajectories diverge as a function of the relational conditions in which the loop operates.

Future research directions follow from these limitations and from the model itself. First, comparative tests across expert domains (e.g., medicine, law, accounting, software engineering, science) would help specify the conditions under which the loop operates similarly and differently. Second, studies that take the relational contingencies as the analytic focus, rather than as background conditions, would let scholars develop a more precise vocabulary for what experts and others build together that shapes how technological change reshapes their work. Third, work on the temporal dynamics of the loop – on how cycles stack, on when frames become sticky, on what kinds of relational moves can revise inferences that have hardened – would extend the theoretical apparatus the dissertation has begun to develop. Fourth, although the dissertation has focused on experts who are navigating technological change, the model's claims about entanglement and recursion should apply equally to settings

where there is no technological interjection. Studies of expert work in mature technological environments would help distinguish the dynamics the model attributes to entangled co-constitution from the dynamics that are specific to periods of technological change.

The dairy farms and veterinary practices in this study are not unusual in the position they occupy. Across professions, the entanglement of knowing, performing, and recognizing has always been there; what is new is that digital technologies are making the entanglement visible and consequential in ways that previous generations of expert work managed to obscure. Audiences see into backstages they once could not see into. Experts adapt their performances to technologies that they did not design and that they cannot fully control. Knowing is reorganized around what can be sensed, logged, and predicted. The three processes were always nested, entangled, and recursive. An ecological ontology offered here is one attempt to keep up. Whether it succeeds in the longer arc of the field's work on expertise and technology will be decided by what the next generation of empirical studies can do with it.

A Closing Note on the Worker

I want to close with a claim that goes beyond the analytic contribution this dissertation has forwarded. It is about the humans whose labor carries our communities, industries, and economies. The biggest privilege in my life has been the opportunity to work with so many kinds of people doing the essential work upon which entire economies are built, and yet who receive so little attention or appreciation. Their labor is largely invisible to those of us who benefit from it – ranchers, fishermen, care givers, janitors, truck drivers, warehouse workers. I have witnessed firsthand the physical, emotional, mental, and spiritual effort that goes into this kind of work, and into most work, which goes unacknowledged in our academic explorations. The complexity and dignity of the human experience get compressed when we are pressured

to objectify our study subjects in pursuit of "theory contribution." Workers and their entanglements with technology become reduced to questions of how to make them more productive when new tools arrive, and workers themselves are exhorted to be "adaptable" in the face of uncertainty in an ambiguous future of work. I have seen how incredibly adaptable workers can be. But demanding that humans continuously bend to structures that those in power refuse to adjust is not a sustainable path forward – not if we wish to keep pursuing technological progress without rendering mentally and spiritually destitute the people on whom the burden of that progress is laid. Humans are adaptable; that is why we have been so successful in the long evolution of life on this planet. But we are also now the architects of our own undoing if we refuse to acknowledge the humanity of our own experience. Technology neither destroys nor creates. Like any tool we clever apes so skillfully manifest, it is part of an engineered ecology of relationships within which we, as social animals, exist and are sustained. The obligation, then, falls on those of us with cultural, social, and economic privilege – those in seats of political influence, in positions of authority within organizations, those who have an audible voice, make decisions, or define the social-psychological timbre of workplaces and communities – to design ways of coexisting with these tools without discarding the humans who constitute the relational scaffolding that gives the tools their meaning. And to try not to dismantle that scaffolding in the process.

The integrative model presented here treats expertise as a relational ecology because the empirical material demanded it. The dairy farmers, veterinarians, breeders, and consultants in this study could not be understood as isolated holders of skill being acted upon by technology. But the choice to study expertise this way is also a choice about what we owe the people whose work we study. Workers are not nodes in production lines. A farmer who has

built a twenty-year relationship with his veterinarian, who gave up his chances at an education to move to a foreign place to blaze a trail for his family, who recognizes each of his cows by sight, is not a person whose expertise reduces to a set of tasks that technology may or may not automate. He is a person whose life is organized around relationships that give the work its meaning, and whose sense of being competent is bound up in being seen, known, and recognized by the others who share those relationships. When technology enters this life, it does not enter at a single point. It reshapes the farmer's relationship to her cows, her peers, her employees, her veterinarian, and herself. The functional changes that most studies of technological change are designed to measure are the smallest part of what is moving.

Research on work, organizations, and technology has tended to shy away from saying this directly. The field's analytic vocabularies – efficiency, augmentation, automation, skill, complementarity, productivity – are powerful, but they keep the worker at arm's length. They treat the worker as the site where technology's effects can be measured rather than as a person whose life is being reorganized by those effects. The distance has its methodological reasons, and our studies are sharper for the discipline it brings. But it comes at a cost. It makes it easier to discuss the future of work in terms of which tasks can be automated, and harder to discuss what kinds of working lives we want people to be able to live. It makes the worker and their lived experience insignificant.

This dissertation has tried, in its modest way, to bring the worker back into focus. The dairy farmers I followed for three years are not the typical experts spotlighted in the technology literature (as one reviewer illustrated pointedly: "Why do we care about dairy farmers?"), and the veterinarians, coworkers, cows, and consultants who shaped their work are not background context. They are the actors whose relationships, evaluations, and accumulated histories

constitute the ecology within which technology has entered and is now reshaping a way of working that has organized lives for generations. If the integrative model contributes anything beyond its analytic moves, I hope it is this: that future research on technology and expert work might take as its starting point not the question of what technology displaces, but the question of what it asks of the people whose lives it enters. That is a harder question. It requires us to look at workers as complex, embedded, relational beings whose work is part of who they are, and to look at technology not as an instrument acting on tasks but as a participant in the relational systems through which workers and their communities sustain meaning over time. The integrity of what we say about workers' lives is the integrity of how we treat them – in our academic examinations and in the organizations we help design.

REFERENCES

Abbott, A. D. (1988). *The system of professions: An essay on the division of expert labor*. University of Chicago Press.

Abdel-Karim, B., Pfeuffer, N., Carl, K. V., & Hinz, O. (2023). How AI-Based Systems Can Induce Reflections: The Case of AI-Augmented Diagnostic Work. *MIS Quarterly, 47*(4), 1395-1424. https://doi.org/10.25300/MISQ/2022/16773

Acemoglu, D., & Restrepo, P. (2019). Artificial Intelligence, Automation, and Work. In *The Economics of Artificial Intelligence* (pp. 197-236). University of Chicago Press.

Alaimo, C., & Kallinikos, J. (2021). Managing by Data: Algorithmic Categories and Organizing. *Organization Studies, 42*(9), 1385-1407. https://doi.org/10.1177/0170840620934062

Alaimo, C., & Kallinikos, J. (2022). Organizations Decentered: Data Objects, Technology and Knowledge. *Organization Science, 33*(1), 19-37. https://doi.org/10.1287/orsc.2021.1552

Allen, R., & Choudhury, P. (Raj). (2022). Algorithm-Augmented Work and Domain Experience: The Countervailing Forces of Ability and Aversion. *Organization Science, 33*(1), 149-169. https://doi.org/10.1287/orsc.2021.1554

Anteby, M. (2013). *Manufacturing Morals: The Values of Silence in Business School Education*. University of Chicago Press.

Anteby, M., Chan, C. K., & DiBenigno, J. (2016). Three Lenses on Occupations and Professions in Organizations: Becoming, Doing, and Relating. *Academy of Management Annals, 10*(1), 183-244. https://doi.org/10.5465/19416520.2016.1120962

Anteby, M., & Holm, A. L. (2021). Translating Expertise across Work Contexts: U.S. Puppeteers Move from Stage to Screen. *American Sociological Review, 86*(2), 310-340. https://doi.org/10.1177/0003122420987199

Anthony, C. (2018). To Question or Accept? How Status Differences Influence Responses to New Epistemic Technologies in Knowledge Work. *Academy of Management Review, 43*(4), 661-679. https://doi.org/10.5465/amr.2016.0334

Anthony, C. (2021). When Knowledge Work and Analytical Technologies Collide: The Practices and Consequences of Black Boxing Algorithmic Technologies. *Administrative Science Quarterly, 66*(4), 1173-1212. https://doi.org/10.1177/00018392211016755

Anthony, C., Bechky, B. A., & Fayard, A.-L. (2023). “Collaborating” with AI: Taking a System View to Explore the Future of Work. *Organization Science, 34*(5), 1672-1694. https://doi.org/10.1287/orsc.2022.1651

Austin, J. R. (2003). Transactive memory in organizational groups: The effects of content, consensus, specialization, and accuracy on group performance. *Journal of Applied Psychology, 88*(5), 866-878. https://doi.org/10.1037/0021-9010.88.5.866

AVMA. (2024). *U.S. veterinarian numbers*. American Veterinary Medical Association. https://www.avma.org/resources-tools/reports-statistics/market-research-statistics-us-veterinarians

Ayyagari, R., Grover, V., & Purvis, R. (2011). Technostress: Technological Antecedents and Implications. *MIS Quarterly, 35*(4), 831-858. https://doi.org/10.2307/41409963

Bailey, D. E., & Barley, S. R. (2020). Beyond design and use: How scholars should study intelligent technologies. *Information and Organization, 30*(2). https://doi.org/10.1016/j.infoandorg.2019.100286

Bailey, D. E., Barley, S. R., & Leonardi, P. M. (2021). Wrestling with Digital Objects and Technologies in Studies of Work. In *Research Methods for Digital Work and Organization* (pp. 25-47). Oxford University Press. https://doi.org/10.1093/oso/9780198860679.003.0002

Bailey, D. E., Faraj, S., Hinds, P. J., Leonardi, P. M., & von Krogh, G. (2022). We Are All Theorists of Technology Now: A Relational Perspective on Emerging Technology and Organizing. *Organization Science, 33*(1), 1-18. https://doi.org/10.1287/orsc.2021.1562

Bailey, D. E., & Leonardi, P. M. (2015). *Technology choices: Why occupations differ in their embrace of new technology*. MIT Press.

Bailey, D. E., Leonardi, P. M., & Barley, S. R. (2012). The Lure of the Virtual. *Organization Science, 23*(5), 1485-1504. https://doi.org/10.1287/orsc.1110.0703

Bar-Haim, Y., Lamy, D., Pergamin, L., Bakermans-Kranenburg, M. J., & Van IJzendoorn, M. H. (2007). Threat-related attentional bias in anxious and nonanxious individuals: A meta-analytic study. *Psychological Bulletin, 133*(1), 1-24. https://doi.org/10.1037/0033-2909.133.1.1

Barley, S. R. (1986). Technology as an occasion for structuring: Evidence from observations of CT scanners and the social order of radiology departments. *Administrative Science Quarterly, 31*(1), 78-108. https://doi.org/10.2307/2392767

Barley, S. R. (1990). The Alignment of Technology and Structure through Roles and Networks. *Administrative Science Quarterly, 35*(1), 61-103. https://doi.org/10.2307/2393551

Barley, S. R. (1996). Technicians in the Workplace: Ethnographic Evidence for Bringing Work into Organizational Studies. *Administrative Science Quarterly, 41*(3), 404-441. https://doi.org/10.2307/2393937

Barley, S. R. (2020). *Work and Technological Change*. Oxford University Press.

Barley, S. R., & Kunda, G. (2001). Bringing Work Back In. *Organization Science, 12*(1), 76-95. https://doi.org/10.1287/orsc.12.1.76.10122

Barrett, M., Oborn, E., Orlikowski, W. J., & Yates, J. (2012). Reconfiguring Boundary Relations: Robotic Innovations in Pharmacy Work. *Organization Science, 23*(5), 1448-1466. https://doi.org/10.1287/orsc.1100.0639

Barton, M. A., & Bunderson, J. S. (2014). Assessing member expertise in groups: An expertise dependence perspective. *Organizational Psychology Review, 4*(3), 228-257. https://doi.org/10.1177/2041386613508975

Baskerville, R. L., Myers, M. D., & Yoo, Y. (2020). Digital First: The Ontological Reversal and New Challenges for Information Systems Research. *MIS Quarterly, 44*(2), 509-524. https://doi.org/10.25300/MISQ/2020/14418

Baudrillard, J. (1994). *Simulacra and Simulation* (S. Glaser, Trans.). University of Michigan Press. https://doi.org/10.3998/mpub.9904

Beane, M. (2019). Shadow Learning: Building Robotic Surgical Skill When Approved Means Fail. *Administrative Science Quarterly, 64*(1), 87-123. https://doi.org/10.1177/0001839217751692

Beane, M., & Anthony, C. (2023). Inverted Apprenticeship: How Senior Occupational Members Develop Practical Expertise and Preserve Their Position When New Technologies Arrive. *Organization Science*, orsc.2023.1688. https://doi.org/10.1287/orsc.2023.1688

Bechky, B. A. (2003). Sharing Meaning Across Occupational Communities: The Transformation of Understanding on a Production Floor. *Organization Science, 14*(3), 312-330. https://doi.org/10.1287/orsc.14.3.312.15162

Bechky, B. A. (2020). Evaluative Spillovers from Technological Change: The Effects of "DNA Envy" on Occupational Practices in Forensic Science. *Administrative Science Quarterly, 65*(3), 606-643. https://doi.org/10.1177/0001839219855329

Bernstein, E. S. (2012). The Transparency Paradox: A Role for Privacy in Organizational Learning and Operational Control. *Administrative Science Quarterly, 57*(2), 181-216. https://doi.org/10.1177/0001839212453028

Bernstein, E. S. (2017). Making Transparency Transparent: The Evolution of Observation in Management Theory. *Academy of Management Annals, 11*(1), 217-266. https://doi.org/10.5465/annals.2014.0076

Boland, R. J., Lyytinen, K., & Yoo, Y. (2007). Wakes of Innovation in Project Networks: The Case of Digital 3-D Representations in Architecture, Engineering, and Construction. *Organization Science, 18*(4), 631-647. https://doi.org/10.1287/orsc.1070.0304

Borgatti, S. P., & Cross, R. (2003). A Relational View of Information Seeking and Learning in Social Networks. *Management Science, 49*(4), 432-445. https://doi.org/10.1287/mnsc.49.4.432.14428

Borgmann, A. (1999). *Holding On to Reality: The Nature of Information at the Turn of the Millennium*. University of Chicago Press. https://doi.org/10.7208/chicago/9780226066226.001.0001

Bossen, C., & Markussen, R. (2010). Infrastructuring and Ordering Devices in Health Care: Medication Plans and Practices on a Hospital Ward. *Computer Supported Cooperative Work (CSCW), 19*(6), 615-637. https://doi.org/10.1007/s10606-010-9131-x

Botelho, T. L. (2024). From Audience to Evaluator: When Visibility into Prior Evaluations Leads to Convergence or Divergence in Subsequent Evaluations Among Professionals. *Organization Science, 35*(5), 1682-1703. https://doi.org/10.1287/orsc.2017.11285

Bourgoin, A., Bencherki, N., & Faraj, S. (2020). "And Who Are You?": A Performative Perspective on Authority in Organizations. *Academy of Management Journal, 63*(4), 1134-1165. https://doi.org/10.5465/amj.2017.1335

Bourgoin, A., & Harvey, J.-F. (2018). Professional image under threat: Dealing with learning-credibility tension. *Human Relations, 71*(12), 1611-1639. https://doi.org/10.1177/0018726718756168

Bourne, L. E., Kole, J. A., & Healy, A. F. (2014). Expertise: Defined, described, explained. *Frontiers in Psychology, 5*, 186. https://doi.org/10.3389/fpsyg.2014.00186

Bowker, G. C., & Star, S. L. (1999). *Sorting things out: Classification and its consequences*. MIT Press.

Boyd, D. (2021). A critical inquiry into the hyperreality of digitalization in construction. *Construction Management and Economics, 39*(7), 549-564. https://doi.org/10.1080/01446193.2021.1904515

Braverman, H. (1998). *Labor and monopoly capital: The degradation of work in the twentieth century* (25th anniversary ed). Monthly Review Press.

Brayne, S. (2017). Big Data Surveillance: The Case of Policing. *American Sociological Review, 82*(5), 977-1008.

Brescoll, V. L. (2011). Who Takes the Floor and Why: Gender, Power, and Volubility in Organizations. *Administrative Science Quarterly, 56*(4), 622-641. https://doi.org/10.1177/0001839212439994

Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.

Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science, 358*(6370), 1530-1534. https://doi.org/10.1126/science.aap8062

Bunderson, J. S. (2003). Recognizing and Utilizing Expertise in Work Groups: A Status Characteristics Perspective. *Administrative Science Quarterly, 48*(4), 557-591. https://doi.org/10.2307/3556637

Bunderson, J. S., & Barton, M. A. (2010). Status cues and expertise assessment in groups: How group members size one another up … and why it matters. In *Status in Management and Organizations* (pp. 215-237). Cambridge University Press. https://doi.org/10.1017/CBO9780511760525.014

Burton-Jones, A., & Grange, C. (2013). From Use to Effective Use: A Representation Theory Perspective. *Information Systems Research, 24*(3), 632-658. https://doi.org/10.1287/isre.1120.0444

Cameron, L. D. (2022). "Making Out" While Driving: Relational and Efficiency Games in the Gig Economy. *Organization Science, 33*(1), 231-252. https://doi.org/10.1287/orsc.2021.1547

Carlile, P. R. (2004). Transferring, Translating, and Transforming: An Integrative Framework for Managing Knowledge Across Boundaries. *Organization Science, 15*(5), 555-568. https://doi.org/10.1287/orsc.1040.0094

Chan, C. K., & Hedden, L. N. (2023). The Role of Discernment and Modulation in Enacting Occupational Values: How Career Advising Professionals Navigate Tensions with Clients. 66(1), 276-305. *Academy of Management Journal*. https://doi.org/10.5465/amj.2020.1014

Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology, 4*(1), 55-81. https://doi.org/10.1016/0010-0285(73)90004-2

Chen, Z., & Chan, J. (2024). Large Language Model in Creative Work: The Role of Collaboration Modality and User Expertise. *Management Science, 70*(12), 9101-9117. https://doi.org/10.1287/mnsc.2023.03014

Christin, A. (2017). Algorithms in practice: Comparing web journalism and criminal justice. *Big Data & Society, 4*(2). https://doi.org/10.1177/2053951717718855

Christin, A. (2020). *Metrics at work: Journalism and the contested meaning of algorithms*. Princeton University Press.

Collins, H. (2010). *Tacit and Explicit Knowledge*. University of Chicago Press.

Collins, H. M., & Evans, R. (2007). *Rethinking expertise*. University of Chicago Press.

Corbin, J. M., & Strauss, A. L. (2015). *Basics of qualitative research: Techniques and procedures for developing grounded theory* (Fourth edition). SAGE.

Cristea, I. C., & Leonardi, P. M. (2019). Get Noticed and Die Trying: Signals, Sacrifice, and the Production of Face Time in Distributed Work. *Organization Science, 30*(3), 552-572. https://doi.org/10.1287/orsc.2018.1265

Dabbish, L., Stuart, C., Tsay, J., & Herbsleb, J. (2012). Social coding in GitHub: Transparency and collaboration in an open software repository. *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work, CSCW '12*, 1277-1286. https://doi.org/10.1145/2145204.2145396

Davis, C. J., & Hufnagel, E. M. (2007). Through the Eyes of Experts: A Socio-Cognitive Perspective on the Automation of Fingerprint Work. *MIS Quarterly, 31*(4), 681-703. https://doi.org/10.2307/25148816

Delfanti, A. (2021). Machinic dispossession and augmented despotism: Digital work in an Amazon warehouse. *New Media & Society, 23*(1), 39-55. https://doi.org/10.1177/1461444819891613

Dell'Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., & Lakhani, K. R. (2023). *Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality*. Harvard Business School Technology & Operations Mgt. Unit Working Paper. https://doi.org/10.2139/ssrn.4573321

DiBenigno, J. (2020). Rapid Relationality: How Peripheral Experts Build a Foundation for Influence with Line Managers. *Administrative Science Quarterly, 65*(1), 20-60. https://doi.org/10.1177/0001839219827006

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err. *Journal of Experimental Psychology: General, 144*(1), 114-126. https://doi.org/10.1037/xge0000033

Dreyfus, H. L., & Dreyfus, S. E. (2005). Peripheral Vision: Expertise in Real World Contexts. *Organization Studies, 26*(5), 779-792. https://doi.org/10.1177/0170840605053102

Edmondson, A. C., Bohmer, R. M., & Pisano, G. P. (2001). Disrupted Routines: Team Learning and New Technology Implementation in Hospitals. *Administrative Science Quarterly, 46*(4), 685-718. https://doi.org/10.2307/3094828

Edmondson, A. C., & Mcmanus, S. E. (2007). Methodological fit in management field research. *Academy of Management Review, 32*(4), 1246-1264. https://doi.org/10.5465/amr.2007.26586086

Endacott, C. G., & Leonardi, P. M. (2022). Artificial Intelligence and Impression Management: Consequences of Autonomous Conversational Agents Communicating on One's Behalf. *Human Communication Research, 48*(3), 462-490. https://doi.org/10.1093/hcr/hqac009

Ericsson, K. A., Hoffman, R. R., Kozbelt, A., & Williams, A. M. (Eds.). (2018). *The Cambridge handbook of expertise and expert performance* (Second edition). Cambridge University Press.

Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review, 100*(3), 363-406. https://doi.org/10.1037/0033-295X.100.3.363

Espeland, W. N., & Stevens, M. L. (1998). Commensuration as a Social Process. *Annual Review of Sociology, 24*, 313-343. https://doi.org/10.1146/annurev.soc.24.1.313

Espeland, W. N., & Stevens, M. L. (2008). A Sociology of Quantification. *European Journal of Sociology, 49*(3), 401-436. https://doi.org/10.1017/S0003975609000150

Evans, J. (2021). How Professionals Construct Moral Authority: Expanding Boundaries of Expert Authority in Stem Cell Science. *Administrative Science Quarterly, 66*(4), 989-1036. https://doi.org/10.1177/00018392211011441

Eyal, G. (2013). For a Sociology of Expertise: The Social Origins of the Autism Epidemic. *American Journal of Sociology, 118*(4), 863-907. https://doi.org/10.1086/668448

Faraj, S., Pachidi, S., & Sayegh, K. (2018). Working and organizing in the age of the learning algorithm. *Information and Organization, 28*(1), 62-70. https://doi.org/10.1016/j.infoandorg.2018.02.005

Faulconbridge, J., Sarwar, A., & Spring, M. (2023). How Professionals Adapt to Artificial Intelligence: The Role of Intertwined Boundary Work. *Journal of Management Studies*, 62(5), 1991-2024. https://doi.org/10.1111/joms.12936

Fayard, A.-L., Gkeredakis, E., & Levina, N. (2016). Framing Innovation Opportunities While Staying Committed to an Organizational Epistemic Stance. *Information Systems Research, 27*(2), 302-323. https://doi.org/10.1287/isre.2016.0623

Fayard, A.-L., Stigliani, I., & Bechky, B. A. (2017). How Nascent Occupations Construct a Mandate: The Case of Service Designers' Ethos. *Administrative Science Quarterly, 62*(2), 270-303. https://doi.org/10.1177/0001839216665805

Floyd, D. L., Prentice-Dunn, S., & Rogers, R. W. (2000). A Meta-Analysis of Research on Protection Motivation Theory. *Journal of Applied Social Psychology, 30*(2), 407-429. https://doi.org/10.1111/j.1559-1816.2000.tb02323.x

Flyverbom, M. (2016). Transparency: Mediation and the Management of Visibilities. *International Journal of Communication, 10*, 110-122.

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change, 114*, 254-280.

Fügener, A., Grahl, J., Gupta, A., & Ketter, W. (2021). Will Humans-in-the-Loop Become Borgs? Merits and Pitfalls of Working with AI. *MIS Quarterly, 45*(3), 1527-1556. https://doi.org/10.25300/MISQ/2021/16553

Galperin, R. V. (2020). Organizational Powers: Contested Innovation and Loss of Professional Jurisdiction in the Case of Retail Medicine. *Organization Science, 31*(2), 508-534. https://doi.org/10.1287/orsc.2019.1314

Geertz, C. (1973). Thick description: Toward an interpretive theory of culture. In *The Interpretation of Cultures* (pp. 3-32). Basic Books.

Gieryn, T. F. (1983). Boundary-Work and the Demarcation of Science from Non-Science: Strains and Interests in Professional Ideologies of Scientists. *American Sociological Review, 48*(6), 781-795. https://doi.org/10.2307/2095325

Glaser, B., & Strauss, A. (1967). *The Discovery of Grounded Theory Strategies for Qualitative Research*. Aldine Transaction.

Glikson, E., & Woolley, A. W. (2020). Human Trust in Artificial Intelligence: Review of Empirical Research. *Academy of Management Annals, 14*(2), 627-660. https://doi.org/10.5465/annals.2018.0057

Gobet, F., & Simon, H. A. (1996). Recall of random and distorted chess positions: Implications for the theory of expertise. *Memory & Cognition, 24*(4), 493-503. https://doi.org/10.3758/BF03200937

Goffman, E. (1956). *The Presentation of Self in Everyday Life*. Doubleday.

Goffman, E. (1967). *Interaction Ritual: Essays in Face-to-Face Behavior*. Routledge. https://doi.org/10.4324/9780203788387

Goodwin, C. (1995). Seeing in Depth. *Social Studies of Science, 25*(2), 237-274. https://doi.org/10.1177/030631295025002002

Gray, J. A. (1982). The neuropsychology of anxiety: An enquiry into the functions of the septo-hippocampal system. *Behavioral and Brain Sciences, 5*(3), 469-484. https://doi.org/10.1017/S0140525X00013066

Greene, J. (2022). *Federal Milk Marketing Orders: An Overview* (No. R45044). Congressional Research Service.

Grønsund, T., & Aanestad, M. (2020). Augmenting the algorithm: Emerging human-in-the-loop work configurations. *The Journal of Strategic Information Systems, 29*(2), 101614. https://doi.org/10.1016/j.jsis.2020.101614

Hadjimichael, D., Ribeiro, R., & Tsoukas, H. (2024). How does Embodiment Enable the Acquisition of Tacit Knowledge in Organizations? From Polanyi to Merleau-Ponty. *Organization Studies*. 45(4). 545-570. https://doi.org/10.1177/01708406241228374

Heimstädt, M., Koljonen, T., & Elmholdt, K. T. (2024). Expertise in Management Research: A Review and Agenda for Future Research. *Academy of Management Annals, 18*(1), 121-156. https://doi.org/10.5465/annals.2022.0078

Heuts, F., & Mol, A. (2013). What Is a Good Tomato? A Case of Valuing in Practice. *Valuation Studies, 1*(2), 125-146. https://doi.org/10.3384/vs.2001-5992.1312125

Hofmann, D. A., Lei, Z., & Grant, A. M. (2009). Seeking help in the shadow of doubt: The sensemaking processes underlying how nurses decide whom to ask for advice. *Journal of Applied Psychology, 94*(5), 1261-1274. https://doi.org/10.1037/a0016557

Hughes, E. C. (1958). *Men and their work*. Free Press.

Huising, R. (2015). To Hive or to Hold? Producing Professional Authority through Scut Work. *Administrative Science Quarterly, 60*(2), 263-299. https://doi.org/10.1177/0001839214560743

International Dairy Foods Association. (2023). *Dairy Delivers: 2023 Economic Impact Study of the U.S. Dairy Products Industry*. IDFA. https://www.idfa.org/news/dairydelivers

Jackson, S. E., & Dutton, J. E. (1988). Discerning Threats and Opportunities. *Administrative Science Quarterly, 33*(3), 370-387. https://doi.org/10.2307/2392714

Jarrahi, M. H. (2019). In the age of the smart artificial intelligence: AI's dual capacities for automating and informating work. *Business Information Review, 36*(4), 178-187. https://doi.org/10.1177/0266382119883999

Jia, N., Luo, X., Fang, Z., & Liao, C. (2024). When and How Artificial Intelligence Augments Employee Creativity. *Academy of Management Journal, 67*(1), 5-32. https://doi.org/10.5465/amj.2022.0426

Jonas, E., McGregor, I., Klackl, J., Agroskin, D., Fritsche, I., Holbrook, C., Nash, K., Proulx, T., & Quirin, M. (2014). Threat and Defense: From Anxiety to Approach. In *Advances in Experimental Social Psychology* (Vol. 49, pp. 219-286). Elsevier. https://doi.org/10.1016/B978-0-12-800052-6.00004-4

Jost, J. T., Glaser, J., Kruglanski, A. W., & Sulloway, F. J. (2003). Political conservatism as motivated social cognition. *Psychological Bulletin, 129*(3), 339-375. https://doi.org/10.1037/0033-2909.129.3.339

Kallinikos, J. (1995). The Architecture of the Invisible: Technology is Representation. *Organization, 2*(1), 117-140. https://doi.org/10.1177/135050849521006

Kallinikos, J. (2009). On the Computational Rendition of Reality: Artefacts and Human Agency. *Organization, 16*(2), 183-202. https://doi.org/10.1177/1350508408100474

Kallinikos, J., Aaltonen, A., & Marton, A. (2010). A theory of digital objects. *First Monday*. https://doi.org/10.5210/fm.v15i6.3033

Kallinikos, J., Aaltonen, A., & Marton, A. (2013). The Ambivalent Ontology of Digital Artifacts. *MIS Quarterly, 37*(2), 357-370.

Kang, S., & Kim, J. W. (2022). The Fragility of Experts: A Moderated-Mediation Model of Expertise, Expert Identity Threat, and Overprecision. *Academy of Management Journal, 65*(2), 577-605. https://doi.org/10.5465/amj.2019.0899

Kaplan, S. (2011). Strategy and PowerPoint: An Inquiry into the Epistemic Culture and Machinery of Strategy Making. *Organization Science, 22*(2), 320-346. https://doi.org/10.1287/orsc.1100.0531

Karunakaran, A. (2022). Status-Authority Asymmetry between Professions: The Case of 911 Dispatchers and Police Officers. *Administrative Science Quarterly, 67*(2), 423-468. https://doi.org/10.1177/00018392211059505

Karunakaran, A., Orlikowski, W. J., & Scott, S. V. (2022). Crowd-Based Accountability: Examining How Social Media Commentary Reconfigures Organizational

Accountability. *Organization Science, 33*(1), 170-193. https://doi.org/10.1287/orsc.2021.1546

Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at Work: The New Contested Terrain of Control. *Academy of Management Annals, 14*(1), 366-410. https://doi.org/10.5465/annals.2018.0174

Kraus, M. W., Park, J. W., & Tan, J. J. X. (2017). Signs of Social Class: The Experience of Economic Inequality in Everyday Life. *Perspectives on Psychological Science, 12*(3), 422-435. https://doi.org/10.1177/1745691616673192

Kronblad, C. (2020). How Digitalization Changes our Understanding of Professional Service Firms. *Academy of Management Discoveries, 6*(3). https://journals.aom.org/doi/abs/10.5465/amd.2019.0027

Lage, C. F. de A., Marques, T. C., Bruno, D. R., Endres, M. I., Ferreira, F., Pires, A. P. A., Leão, K., & de Lima, F. S. (2024). Farmers' Perceptions on Implementing Automatic Milking Systems in Large USA Dairies: Decision-Making Process, Management Practices, Labor, and Herd Performance. *Animals : An Open Access Journal from MDPI, 14*(2), 218. https://doi.org/10.3390/ani14020218

Latour, B., & Woolgar, S. (1986). *Laboratory life: The construction of scientific facts*. Princeton University Press.

Lazarus, R. S., & Folkman, S. (1984). *Stress, appraisal, and coping* (11. [print.]). Springer.

Leavell, V. (2023). The Future Isn't What It Used to Be: Anticipatory Control and Digital Change in Water Infrastructure. *Academy of Management Proceedings, 2023*(1), 12941. https://doi.org/10.5465/AMPROC.2023.12941abstract

Lebovitz, S., Levina, N., & Lifshitz-Assaf, H. (2021). Is AI Ground Truth Really True? The Dangers of Training and Evaluating AI Tools Based on Experts' Know-What. *MIS Quarterly, 45*, 1501-1526. https://doi.org/10.25300/MISQ/2021/16564

Lebovitz, S., Lifshitz-Assaf, H., & Levina, N. (2022). To Engage or Not to Engage with AI for Critical Judgments: How Professionals Deal with Opacity When Using AI for Medical Diagnosis. *Organization Science, 33*(1), 126-148. https://doi.org/10.1287/orsc.2021.1549

Lei, Y.-W., & Kim, R. (2024). Automation and Augmentation: Artificial Intelligence, Robots, and Work. *Annual Review of Sociology, 50*(1), 251-272. https://doi.org/10.1146/annurev-soc-090523-050708

Leonardi. (2011). When Flexible Routines Meet Flexible Technologies: Affordance, Constraint, and the Imbrication of Human and Material Agencies. *MIS Quarterly, 35*(1), 147-167. https://doi.org/10.2307/23043493

Leonardi, P. M. (2007). Activating the Informational Capabilities of Information Technology for Organizational Change. *Organization Science, 18*(5), 813-831. https://doi.org/10.1287/orsc.1070.0284

Leonardi, P. M. (2009). Why Do People Reject New Technologies and Stymie Organizational Changes of Which They Are in Favor? Exploring Misalignments Between Social Interactions and Materiality. *Human Communication Research, 35*(3), 407-441. https://doi.org/10.1111/j.1468-2958.2009.01357.x

Leonardi, P. M. (2012). *Car Crashes Without Cars: Lessons about Simulation Technology and Organizational Change from Automotive Design*. MIT Press.

Leonardi, P. M. (2013). When Does Technology Use Enable Network Change in Organizations? A Comparative Study of Feature Use and Shared Affordances. *MIS Quarterly, 37*(3), 749-775. https://doi.org/10.25300/MISQ/2013/37.3.04

Leonardi, P. M. (2014). Social Media, Knowledge Sharing, and Innovation: Toward a Theory of Communication Visibility. *Information Systems Research, 25*(4), 796-816. https://doi.org/10.1287/isre.2014.0536

Leonardi, P. M. (2015). Ambient Awareness and Knowledge Acquisition: Using Social Media to Learn "Who Knows What" and "Who Knows Whom." *MIS Quarterly, 39*(4), 747-762.

Leonardi, P. M. (2018). Social Media and the Development of Shared Cognition: The Roles of Network Expansion, Content Integration, and Triggered Recalling. *Organization Science, 29*(4), 547-568. https://doi.org/10.1287/orsc.2017.1200

Leonardi, P. M., & Barley, S. R. (2010). What's Under Construction Here? Social Action, Materiality, and Power in Constructivist Studies of Technology and Organizing. *Academy of Management Annals, 4*(1), 1-51. https://doi.org/10.5465/19416521003654160

Leonardi, P. M., & Treem, J. W. (2012). Knowledge management technology as a stage for strategic self-presentation: Implications for knowledge sharing in organizations. *Information and Organization, 22*(1), 37-59. https://doi.org/10.1016/j.infoandorg.2011.10.003

Leonardi, P. M., & Treem, J. W. (2020). Behavioral Visibility: A new paradigm for organization studies in the age of digitization, digitalization, and datafication. *Organization Studies, 41*(12), 1601-1625. https://doi.org/10.1177/0170840620970728

Lewis, K. (2004). Knowledge and Performance in Knowledge-Worker Teams: A Longitudinal Study of Transactive Memory Systems. *Management Science, 50*(11), 1519-1533. https://doi.org/10.1287/mnsc.1040.0257

Libby, R., Trotman, K. T., & Zimmer, I. (1987). Member variation, recognition of expertise, and group performance. *Journal of Applied Psychology, 72*(1), 81-87. https://doi.org/10.1037/0021-9010.72.1.81

Lifshitz-Assaf, H. (2018). Dismantling Knowledge Boundaries at NASA: The Critical Role of Professional Identity in Open Innovation. *Administrative Science Quarterly, 63*(4), 746-782. https://doi.org/10.1177/0001839217747876

Lindebaum, D., Vesa, M., & Den Hond, F. (2020). Insights From "The Machine Stops" to Better Understand Rational Assumptions in Algorithmic Decision Making and Its Implications for Organizations. *Academy of Management Review, 45*(1), 247-263. https://doi.org/10.5465/amr.2018.0181

Littlepage, G., Robison, W., & Reddington, K. (1997). Effects of Task Experience and Group Experience on Group Performance, Member Ability, and Recognition of Expertise. *Organizational Behavior and Human Decision Processes, 69*(2), 133-147. https://doi.org/10.1006/obhd.1997.2677

Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes, 151*, 90-103. https://doi.org/10.1016/j.obhdp.2018.12.005

Lynch, M. (1991). Science in the age of mechanical reproduction: Moral and epistemic relations between diagrams and photographs. *Biology & Philosophy, 6*(2), 205-226. https://doi.org/10.1007/BF02426838

Marlow, J., Dabbish, L., & Herbsleb, J. (2013). Impression formation in online peer production: Activity traces and personal profiles in GitHub. *Proceedings of the 2013 Conference on Computer Supported Cooperative Work, CSCW '13*, 117-128. https://doi.org/10.1145/2441776.2441792

Marti, E., & Gond, J.-P. (2018). When Do Theories Become self-Fulfilling? Exploring the Boundary Conditions of Performativity. *Academy of Management Review, 43*(3), 487-508. https://doi.org/10.5465/amr.2016.0071

Mazmanian, M. (2013). Avoiding the Trap of Constant Connectivity: When Congruent Frames Allow for Heterogeneous Practices. *Academy of Management Journal, 56*(5), 1225-1250. https://doi.org/10.5465/amj.2010.0787

Mazmanian, M., & Beckman, C. M. (2018). "Making" Your Numbers: Engendering Organizational Control Through a Ritual of Quantification. *Organization Science, 29*(3), 357-379. https://doi.org/10.1287/orsc.2017.1185

Mol, A. (2000). What Diagnostic Devices Do: The Case of Blood Sugar Measurement. *Theoretical Medicine and Bioethics, 21*(1), 9-22. https://doi.org/10.1023/A:1009999119586

Monteiro, E., & Parmiggiani, E. (2019). Synthetic Knowing: The Politics of the Internet of Things. *MIS Quarterly, 43*(1), 167-184. https://doi.org/10.25300/MISQ/2019/13799

Moreland, R. L., & Myaskovsky, L. (2000). Exploring the Performance Benefits of Group Training: Transactive Memory or Improved Communication? *Organizational Behavior and Human Decision Processes, 82*(1), 117-133. https://doi.org/10.1006/obhd.2000.2891

Muniesa, F. (2011). A Flank Movement in the Understanding of Valuation. *The Sociological Review, 59*(2_suppl), 24-38. https://doi.org/10.1111/j.1467-954X.2012.02056.x

Nardi, B. A., & Engeström, Y. (1999). A Web on the Wind: The Structure of Invisible Work. *Computer Supported Cooperative Work: The Journal of Collaborative Computing, 8*, 1-8. https://doi.org/10.1023/A:1008694621289

Nelsen, B. J., & Barley, S. R. (1997). For Love or Money? Commodification and the Construction of an Occupational Mandate. *Administrative Science Quarterly, 42*(4), 619-653. https://doi.org/10.2307/2393652

Nelson, A., Anthony, C., & Tripsas, M. (2023). "If I Could Turn Back Time": Occupational Dynamics, Technology Trajectories, and the Reemergence of the Analog Music Synthesizer. *Administrative Science Quarterly, 68*(2), 551-599. https://doi.org/10.1177/00018392231163178

Nelson, A. J., & Irwin, J. (2014). "Defining What We Do–All Over Again": Occupational Identity, Technological Change, and the Librarian/Internet-Search Relationship. *Academy of Management Journal, 57*(3), 892-928. https://doi.org/10.5465/amj.2012.0201

Nonaka, I., & Von Krogh, G. (2009). Perspective—Tacit Knowledge and Knowledge Conversion: Controversy and Advancement in Organizational Knowledge Creation Theory. *Organization Science, 20*(3), 635-652. https://doi.org/10.1287/orsc.1080.0412

Orlikowski, W. J. (1996). Improvising Organizational Transformation Over Time: A Situated Change Perspective. *Information Systems Research, 7*(1), 63-92.

Orlikowski, W. J. (2000). Using Technology and Constituting Structures: A Practice Lens for Studying Technology in Organizations. *Organization Science, 11*(4), 404-428. https://doi.org/10.1287/orsc.11.4.404.14600

Orlikowski, W. J. (2002). Knowing in Practice: Enacting a Collective Capability in Distributed Organizing. *Organization Science, 13*(3), 249-273.

Orlikowski, W. J., & Gash, D. C. (1994). Technological frames: Making sense of information technology in organizations. *ACM Transactions on Information Systems, 12*(2), 174-207. https://doi.org/10.1145/196734.196745

Orlikowski, W. J., & Scott, S. V. (2008). Sociomateriality: Challenging the Separation of Technology, Work and Organization. *Academy of Management Annals, 2*(1), 433-474. https://doi.org/10.5465/19416520802211644

Orlikowski, W. J., & Scott, S. V. (2014). What Happens When Evaluation Goes Online? Exploring Apparatuses of Valuation in the Travel Sector. *Organization Science, 25*(3), 868-891. https://doi.org/10.1287/orsc.2013.0877

Østerlie, T., & Monteiro, E. (2020). Digital sand: The becoming of digital representations. *Information and Organization, 30*(1), 100275. https://doi.org/10.1016/j.infoandorg.2019.100275

Pachidi, S., Berends, H., Faraj, S., & Huysman, M. (2021). Make Way for the Algorithms: Symbolic Actions and Change in a Regime of Knowing. *Organization Science, 32*(1), 18-41. https://doi.org/10.1287/orsc.2020.1377

Pakarinen, P., & Huising, R. (2023). Relational Expertise: What Machines Can't Know. *Journal of Management Studies*, joms.12915. https://doi.org/10.1111/joms.12915

Patriotta, G., & Spedale, S. (2009). Making Sense Through Face: Identity and Social Interaction in a Consultancy Task Force. *Organization Studies, 30*(11), 1227-1248. https://doi.org/10.1177/0170840609347036

Petriglieri, J. L. (2011). Under Threat: Responses to and the Consequences of Threats to Individuals' Identities. *Academy of Management Review, 36*(4), 641-662. https://doi.org/10.5465/amr.2009.0087

Pinch, T. (1996). The Social Construction of Technology: A Review. In *Technological Change*. Routledge.

Pinch, T. (2010). The Invisible Technologies of Goffman's Sociology From the Merry-Go-Round to the Internet. *Technology and Culture, 51*(2), 409-424. https://doi.org/10.1353/tech.0.0456

Polanyi, M. (2009). *The Tacit Dimension* (A. Sen, Ed.). University of Chicago Press.

Porter, T. M. (1995). *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton University Press.

Prasad, P. (1993). Symbolic Processes in the Implementation of Technological Change: A Symbolic Interactionist Study of Work Computerization. *Academy of Management Journal, 36*(6), 1400-1429. https://doi.org/10.5465/256817

Pratt, M. G., Rockmann, K. W., & Kaufmann, J. B. (2006). Constructing Professional Identity: The Role of Work and Identity Learning Cycles in the Customization of Identity Among Medical Residents. *Academy of Management Journal, 49*(2), 235-262. https://doi.org/10.5465/amj.2006.20786060

Raisch, S., & Krakowski, S. (2021). Artificial Intelligence and Management: The Automation-Augmentation Paradox. *Academy of Management Review, 46*(1), 192-210. https://doi.org/10.5465/amr.2018.0072

Recker, J., Lukyanenko, R., Jabbari, M., Samuel, B. M., & Castellanos, A. (2021). From Representation to Mediation: A New Agenda for Conceptual Modeling Research in a Digital World. *MIS Quarterly, 45*(1), 269-300. https://doi.org/10.25300/MISQ/2021/16027

Ren, Y., & Argote, L. (2011). Transactive Memory Systems 1985-2010: An Integrative Framework of Key Dimensions, Antecedents, and Consequences. *Academy of Management Annals, 5*(1), 189-229. https://doi.org/10.5465/19416520.2011.590300

Robey, D. (1979). User Attitudes and Management Information System. *Academy of Management Journal, 22*(3), 527-538.

Rousseau, D. M., & Stouten, J. (2025). Experts and Expertise in Organizations: An Integrative Review on Individual Expertise. *Annual Review of Organizational Psychology and Organizational Behavior, 12*, 159-184.

Sandberg, J., Rouleau, L., Langley, A., & Tsoukas, H. (2017). *Skillful Performance: Enacting Capabilities, Knowledge, Competence, and Expertise in Organizations*. Oxford University Press.

Scarbrough, H., Chen, Y., & Patriotta, G. (2024). The AI of the Beholder: Intra-Professional Sensemaking of an Epistemic Technology. *Journal of Management Studies*, joms.13065. https://doi.org/10.1111/joms.13065

Schmitz, A. (2025, November 5). *2024 U.S. dairy statistics: Annual milk production declined again*. Progressive Dairy. Ag Proud. https://www.agproud.com/articles/62420-2024-us-dairy-statistics-annual-milk-production-declined-again

Schnitkey, C. Z., and G. (2024). US Dairy Herds and Policy and the 2022 Census of Agriculture. *Farmdoc Daily, 14*(38). https://farmdocdaily.illinois.edu/2024/02/us-dairy-herds-and-policy-and-the-2022-census-of-agriculture.html

Sergeeva, A. V., Faraj, S., & Huysman, M. (2020). Losing Touch: An Embodiment Perspective on Coordination in Robotic Surgery. *Organization Science, 31*(5), 1248-1271. https://doi.org/10.1287/orsc.2019.1343

Shami, N. S., Ehrlich, K., Gay, G., & Hancock, J. T. (2009). Making sense of strangers' expertise from signals in digital artifacts. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 69-78. https://doi.org/10.1145/1518701.1518713

Shih, M., Pittinsky, T. L., & Ambady, N. (1999). Stereotype susceptibility: Identity salience and shifts in quantitative performance. *Psychological Science, 10*(1), 80-83.

Simon, H. A. (1972). Theories of bounded rationality. *Decision and Organization, 1*(1), 161-176.

Spencer, S. J., Logel, C., & Davies, P. G. (2016). Stereotype Threat. *Annual Review of Psychology, 67*(1), 415-437. https://doi.org/10.1146/annurev-psych-073115-103235

Spradley, J. P. (1979). *The Ethnographic Interview*. Holt, Rinehart, and Winston.

Stanaland, A., Gaither, S., & Gassman-Pines, A. (2023). When Is Masculinity "Fragile"? An Expectancy-Discrepancy-Threat Model of Masculine Identity. *Personality and Social Psychology Review, 27*(4), 359-377.

Star, S. L., & Strauss, A. (1999). Layers of Silence, Arenas of Voice: The Ecology of Visible and Invisible Work. *Computer Supported Cooperative Work (CSCW), 8*(1), 9-30. https://doi.org/10.1023/A:1008651105359

Staw, B. M., Sandelands, L. E., & Dutton, J. E. (1981). Threat Rigidity Effects in Organizational Behavior: A Multilevel Analysis. *Administrative Science Quarterly, 26*(4), 501-524. https://doi.org/10.2307/2392337

Steele, C. M. (1997). A threat in the air: How Stereotypes Shape Intellectual Identity and Performance. *American Psychologist, 52*(6), 613-629.

Stohl, C., Stohl, M., & Leonardi, P. M. (2016). Managing Opacity: Information Visibility and the Paradox of Transparency in the Digital Age. *International Journal of Communication, 10*, 123-137.

Strauss, A. (1978). A social world perspective. *Studies in Symbolic Interaction, 1*(1), 119-128.

Su, C. (2012). Who Knows Who Knows What in the Group? The Effects of Communication Network Centralities, Use of Digital Knowledge Repositories, and Work Remoteness on Organizational Members' Accuracy in Expertise Recognition. *Communication Research, 39*(4), 614-640.

Suddaby, R., & Greenwood, R. (2005). Rhetorical Strategies of Legitimacy. *Administrative Science Quarterly, 50*(1), 35-67. https://doi.org/10.2189/asqu.2005.50.1.35

Treem, J. W. (2012). Communicating Expertise: Knowledge Performances in Professional-Service Firms. *Communication Monographs, 79*(1), 23-47. https://doi.org/10.1080/03637751.2011.646487

Treem, J. W. (2013). Technology Use as a Status Cue: The Influences of Mundane and Novel Technologies on Knowledge Assessments in Organizations. *Journal of Communication, 63*(6), 1032-1053. https://doi.org/10.1111/jcom.12061

Treem, J. W. (2016). How Organizations Communicate Expertise Without Experts: Practices and Performances of Knowledge-Intensive Firms. *Management Communication Quarterly, 30*(4), 503-531. https://doi.org/10.1177/0893318916635750

Treem, J. W., Dailey, S. L., Pierce, C. S., & Leonardi, P. M. (2015). Bringing Technological Frames to Work: How Previous Experience with Social Media Shapes the Technology's Meaning in an Organization. *Journal of Communication, 65*(2), 396-422. https://doi.org/10.1111/jcom.12149

Treem, J. W., & Leonardi, P. M. (2017). Recognizing Expertise: Factors Promoting Congruity Between Individuals' Perceptions of Their Own Expertise and the Perceptions of Their Coworkers. *Communication Research, 44*(2), 198-224. https://doi.org/10.1177/0093650215605154

Truelove, E., & Kellogg, K. C. (2016). The Radical Flank Effect and Cross-occupational Collaboration for Technology Development during a Power Shift. *Administrative Science Quarterly, 61*(4), 662-701. https://doi.org/10.1177/0001839216647679

UC Davis. (2020). *Veterinarian Workforce and Education in California*. University of California, Davis.

USDA Agricultural Marketing Service. (2023). *Federal Milk Marketing Orders*. https://www.ams.usda.gov/rules-regulations/moa/dairy

USDA National Agricultural Statistics Service. (2025). *Milk Production 02/21/2025*. https://www.nass.usda.gov/Publications/Todays_Reports/reports/mkpr0225.pdf

Vaast, E., & Pinsonneault, A. (2021). When Digital Technologies Enable and Threaten Occupational Identity: The Delicate Balancing Act of Data Scientists. *MIS Quarterly, 45*(3), 1087-1112. https://doi.org/10.25300/MISQ/2021/16024

Valentine, M., & Hinds, R. (2022). How Algorithms Change Occupational Expertise by Prompting Explicit Articulation and Testing of Experts' Theories. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4246167

Vallas, S. P. (2001). Symbolic boundaries and the new division of labor: Engineers, workers and the restructuring of factory life. *Research in Social Stratification and Mobility, 18*, 3-37. https://doi.org/10.1016/S0276-5624(01)80021-4

Van Den Broek, E., Sergeeva, A., & Huysman, M. (2021). When the Machine Meets the Expert: An Ethnography of Developing AI for Hiring. *MIS Quarterly, 45*(3), 1557-1580. https://doi.org/10.25300/MISQ/2021/16559

Van Doorn, S., Georgakakis, D., Oehmichen, J., & Reimer, M. (2023). Opportunity or Threat? Exploring Middle Manager Roles in the Face of Digital Transformation. *Journal of Management Studies, 60*(7), 1684-1719. https://doi.org/10.1111/joms.12880

Van Maanen, J. (2011). *Tales of the Field: On Writing Ethnography* (2nd ed). University of Chicago Press.

von Krogh, G. (2018). Artificial Intelligence in Organizations: New Opportunities for Phenomenon-Based Theorizing. *Academy of Management Discoveries, 4*(4), 404-409. https://doi.org/10.5465/amd.2018.0084

Waardenburg, L., Huysman, M., & Sergeeva, A. V. (2022). In the Land of the Blind, the One-Eyed Man Is King: Knowledge Brokerage in the Age of Learning Algorithms. *Organization Science, 33*(1), 59-82. https://doi.org/10.1287/orsc.2021.1544

Walton, G. M., Murphy, M. C., & Ryan, A. M. (2015). Stereotype threat in organizations: Implications for equity and performance. *Annual Review of Organizational Psychology and Organizational Behavior, 2*, 523-550.

Wang, W., Gao, G. (Gordon), & Agarwal, R. (2024). Friend or Foe? Teaming Between Artificial Intelligence and Workers with Variation in Experience. *Management Science, 70*(9), 5753-5775. https://doi.org/10.1287/mnsc.2021.00588

Weber, R. (2003). Editor's Comments: Still Desperately Seeking the IT Artifact. *MIS Quarterly, 27*(3), iii-xi.

Weick, K. E. (1990). Technology as equivoque: Sensemaking in new technologies. In *Technology and organizations* (pp. 1-44). Jossey-Bass.

Weick, K. E. (1995). *Sensemaking in organizations*. Sage Publications.

Whitt, C. (2024). *Overview of Dairy Cooperatives* (No. R48285). Congressional Research Service.

Wu, L., & Kane, G. C. (2021). Network-Biased Technical Change: How Modern Digital Collaboration Tools Overcome Some Biases but Exacerbate Others. *Organization Science, 32*(2), 273-292. https://doi.org/10.1287/orsc.2020.1368

Yam, K. C., Tang, P. M., Jackson, J. C., Su, R., & Gray, K. (2023). The rise of robots increases job insecurity and maladaptive workplace behaviors: Multimethod evidence. *Journal of Applied Psychology, 108*(5), 850-870. https://doi.org/10.1037/apl0001045

Zuboff, S. (1988). *In the age of the smart machine: The future of work and power*. Basic Books.

Zuboff, S. (2015). Big other: Surveillance Capitalism and the Prospects of an Information Civilization. *Journal of Information Technology, 30*(1), 75-89. https://doi.org/10.1057/jit.2015.5

APPENDIX A: Conventional Versus Data-intensive Dairy Operations

The table below summarizes the key technologies, techniques, and processes that differed between conventional and data-intensive dairy farm operations across the three studies in this dissertation. "Conventional" and "data-intensive" function as analytic poles on a continuum; in practice, most farms occupied positions along this spectrum, with varying combinations of these conventional and technological methods in use.

Activity	Conventional approach	Data-intensive approach
I. Identifying and tracking individual cows		
Cow identification	Visual ID via ear tags, colored ankle bands, chalk/paint marks on cow's bodies, and hand-written notes on paper; cows are identified and sorted through visual identification.	RFID tags are read automatically by handheld wand scanners, sort gates, rotary stalls, and robotic milkers; cows are identified and sorted using sensor-enabled devices.
II. Milking and milk data capture		
Milk extraction	Cows are moved to a central milking parlor in pen groups 2-3 times per day; milk is pooled in bulk tanks and volume is measured per tank.	Robotic milkers milk cows individually 2-4 times per day; sensor-equipped rotaries milk cows in pen groups 2-3 times per day; per-cow, per-quarter data of yield, flow rate, milk quality, and volume are recorded in real time.
Milk quality testing (e.g., conductivity, fat, protein, and somatic cell count)	Milk is tested at processing plants after delivery; results are emailed to farms hours or days later; one measurement per metric per tank which could include milk from hundreds of cows.	In-line milk analyzers test each cow's milk in real time by udder quarter; results are rendered within seconds and automatically integrated into the cow's record.
Milk records	Pen-level averages are estimated and entered manually into herd management software to calculate average milk production per cow; no individualized data for milk quality.	Continuous, longitudinal data on milk production, milk quality, and trends are automatically logged into each individual cow's record and can be accessed from any device.
III. Heat detection and breeding		
Heat detection	A single experienced breeder walks cow pens daily, visually	Activity-tracking collars with accelerometry sensors capture data

Activity	Conventional approach	Data-intensive approach
	looking for symptoms or proxies of ovulation (e.g., vaginal discharge, rubbed-off tail chalk, active mounting behavior).	on movement, rumination, feeding, and resting 24/7.
Triggers for artificial insemination	Breeders' embodied judgment integrates proxy clues to assess heat likelihood; manual rectal palpation confirms uterine "tone."	Algorithm computes heat likelihood and optimal breeding window per cow; pre-determined algorithmic threshold triggers cow inclusion in breeding lists.
Insemination and recording	Breeder inseminates on the spot, often 2-3 days in a row; breeding events are entered manually into herd management software in office.	Heat events sync into herd management software; breeder inseminates based on computer-generated list and confirms events in phone app that sync back into herd management software.
IV. Mastitis and udder health		
Mastitis detection	Milkers and herdsmen monitor for symptoms by visual and tactile examination during milking or routine herd checks (e.g., palpating udder for inflammation or swelling, observing milk for clots or abnormality).	Robots and algorithms detect chemical and physical signatures calculated from milk conductivity changes, somatic cell count increases, and milk yield drops per quarter; algorithm automatically flags and separate cows predicted to have mastitis for examination and treatment.
Timing	Identified at clinical stage when symptoms are visible to humans.	Subclinical detection 1-2 days before symptoms appear.
Treatment	Treatment decisions are based on the farmer's hands-on assessment of physical symptoms and rudimentary tests, such as milk culture plates that take 1 – 2 days to incubate.	"Udder Health" reports show at-risk cows with color-coded probability scores; "Mastitis likelihood" scores generate treatment lists.
V. Illness and general health monitoring		
Detection	Herdsmen and milkers visually monitor cows during routine work; intuitive "hunches" about cow comfort (e.g., ear position, posture,	Activity collars and robotic milkers detect deviations from each cow's individual baseline (e.g., reduced rumination, decreased movement,

Activity	Conventional approach	Data-intensive approach
	eyes, body language, breathing/coughing, nasal discharge).	abnormal resting, increased milk temperature, decreased milk production).
Decision support	Sick cows are identified primarily through embodied pattern recognition built over years of working directly with cows and developing "cow sense."	"Sick chance" algorithmic scores are generated from combined sensor inputs from activity tracking collars and robotic milkers; cows flagged on "Attention Cows" list for farmers to examine and treat.
VI. Breeding strategy and genetics		
Genetic evaluation	Bull selection based on phenotypic observation (e.g., body size and conformation, udder shape and size, health, milk production, etc.), lineage records, and breeder/farmer preference.	Genomic testing (DNA analysis) predicts heritability across 30+ traits; composite merit indexes (e.g., Dairy Wellness Profit Index) rank animals' productivity potential early in life.
Breeding strategy	Breeding decisions are made cow-by-cow based on farmer knowledge and episodic evaluations; bulls are selected based on qualitative assessments of cow and bull traits to produce desired traits; outcomes are realized when calves grow to production age.	Genomic service providers match cows with bulls depending on desired traits; algorithmic rankings stratify herd – top percentiles are bred to sexed semen for replacement milk calves; bottom percentile are bred to beef (Angus) bulls to terminate genetic lines; outcomes are predicted by models before cows ever reach reproductive age.
VII. Pregnancy confirmation		
Diagnosis	Veterinarians conduct rectal palpation and handheld ultrasound of fetus to confirm pregnancies and diagnose fetal conditions; routine confirmations are conducted at scheduled intervals throughout gestation.	Blood tests measure proteins indicating early pregnancy; threshold-based diagnosis is configurable in herd management software; results auto-populate into cow records and veterinarians conduct follow ups.
VIII. Culling decisions		
Decision basis	Decisions are based on direct observation of individual cow's	Algorithmic "cull lists" are generated from projected

Activity	Conventional approach	Data-intensive approach
	body condition, recall of their history, available performance metrics, and farmer/manager evaluation of a cow's "goodness."	performance based on data records and trends; comparative rankings position individual cows along a herd-wide distribution (e.g., relative value scores).
IX. Record keeping and herd management		
Volume of data	Sparse, discrete entries – one cow's record typically includes around 30 discrete events across a cow's career.	Comprehensive, continuous digital representation – one cow's record could include thousands of data points, modeled predictions, labels, diagnoses, and discrete events.
Data entry	Data is entered manually (e.g., treatments, breedings, hoof trims, pregnancy checks) into one centralized herd management software program with many events not recorded at all.	Mix of auto-captured sensor data and manual entries; data is distributed across multiple platforms and integrated via APIs or digital boundary objects (e.g., DairyComp, DelPro, Horizon, Excel spreadsheets, web platforms).
Reporting	Reports are generated periodically as discrete snapshots and reviewed in batches.	Software programs produce real-time dashboards, color-coded data tables, push alerts, and daily action lists that are updated continuously.
X. External specialist involvement		
Access to farm activity/data	Veterinarians, breeders, nutritionists, hoof trimmers, and other services providers visit on regular schedules; they provide specialized expertise and rely on farmer accounts for information and on – site observation during visits to make recommendations.	Same specialists, but with continuous remote access to digital records anytime and from anywhere; vets and other service providers can access farm records, interrogate data, benchmark across farms, and identify discrepancies independently of farmer mediation.

APPENDIX B: Table of Field Sites

Site ID	Site Type	Herd Size	Employees	Milking System	Key Technologies in Use
CR1	Calf ranch	~2,000	10	N/A	Ear tag identification, herd management software
CR2	Calf ranch	~4,000	6	N/A	Ear tag identification, herd management software
DF1	Dairy farm	~6,000	>30	Rotary	Rotary parlor, herd management software
DF2	Dairy farm	~5,000	>30	Hybrid (Lely)	Robotic milking, activity collars, herd management software
DF6	Dairy farm	~5,400	>30	Hybrid (DeLaval)	Robotic milking, parallel parlor, activity collars (partial herd), ear tags, genomics, herd management software
DF9	Dairy farm	~2,000	22	Robotic (Lely)	Robotic milking, activity collars, genomics, herd management software
DF10	Dairy farm	~1,200	11	Robotic (Lely)	Robotic milking, activity collars, genomics, herd management software
DF11	Dairy farm	~480	7	Robotic (Lely)	Robotic milking, activity collars, genomics, herd management software
DF12	Dairy farm	~2100	–	Hybrid (DeLaval)	Robotic milking, parallel parlor, activity collars, herd management software
DF13	Dairy farm	~1,600	–	Conventional	Parallel parlor, activity collars, herd management software
DF14	Dairy farm	~1,800	–	Conventional	Parallel parlor, activity collars, genomics, herd management software
DF15	Dairy farm	~1,200	–	Conventional	Parallel parlor, herd management software
DF16	Dairy farm	~1400	–	Hybrid (DeLaval)	Robotic milking, rotary parlor, activity collars, herd management software
DF17	Dairy farm	~6,700	22	Conventional	Parallel parlor, activity collars, herd management software
DF18	Dairy farm	~6,000	>30	Conventional	Parallel parlor, activity collars, herd management software

Site ID	Site Type	Herd Size	Employees	Milking System	Key Technologies in Use
DF19	Dairy farm	~6,500	>30	Rotary	Rotary parlor, activity collars, herd management software
DF20	Dairy farm	~10,000	>30	Rotary	Rotary parlor, activity collars, herd management software
DF21	Dairy farm	~1,000	–	Hybrid (DeLaval)	Robotic milking, rotary parlor, activity collars, genomics, herd management software

Notes. Herd size figures are approximations reported by farm owners or managers at the time of fieldwork. “Hybrid” denotes farms operating both robotic and conventional milking systems concurrently. “Robotic milking” denotes the use of automated robotic milking systems. DeLaval and Lely are the two main companies providing robotic milking systems and associated services. “Parallel parlors” are typically conventional but can sometimes be retrofitted with sensors. “Rotary parlors” are equipped with per-stall sensors. “Activity collars” refers to sensor-equipped neck collars that track cow movement, rumination, and reproductive cycles. “Genomics” indicates use of genomic testing for breeding and herd management decisions. “Herd management software” refers to DairyComp, DHI-Plus, or other legacy databases used on all farms to record and manage digital records. All farms used RFID tags. All site identifiers are pseudonymous.

APPENDIX C: Table of Informants

ID	Site	Pseudonym	Organization	Role	Tenure (yrs)
F1	DF1	Ben	Dairy farm	Farm owner	50+
F2	DF1	Bradley	Dairy farm	Farm manager	10+
F36	DF2	Fiona	Dairy farm	Farm owner / manager	40+
F38	DF2	Frank	Dairy farm	Robot barn manager	4
F55	DF6	Derek	Dairy farm	Farm owner / manager	20+
F56	DF6	David	Dairy farm	Farm owner / assistant manager	3
F57	DF6	Daniel	Dairy farm	Robot barn manager	20+
F58	DF6	–	Dairy farm	Herdsman	–
F59	DF6	Devon	Dairy farm	Head herdsman	–
F60	DF6	Damien	Dairy farm	Manager (hospital)	20+
F61	DF6	Dylan	Dairy farm	Herd manager	23
F62	DF6	–	Dairy farm	Milker	–
F63	DF6	–	Dairy farm	Milker	–
F66	DF6	–	Dairy farm	Calf feeder	–
F67	DF6	–	Dairy farm	Herdsman	–
F68	DF6	–	Dairy farm	Herdsman	–
F70	DF6	–	Dairy farm	Herdsman	–
F71	DF6	–	Dairy farm	Milker	–
F74	DF6	–	Dairy farm	Herdsman (calves)	–
F75	DF6	–	Dairy farm	Herdsman	–
F79	DF6	–	Dairy farm	Herdsman (hospital)	–
F80	DF6	–	Dairy farm	Herdsman	–
F81	DF6	–	Dairy farm	Herdsman	–
F82	DF6	–	Dairy farm	Manager (milking parlor)	20+
F99	DF9	Carl	Dairy farm	Farm owner	10+
F100	DF9	Chris	Dairy farm	Farm manager	20+
F101	DF9	–	Dairy farm	Herdsman	–

ID	Site	Pseudonym	Organization	Role	Tenure (yrs)
F102	DF9	Cameron	Dairy farm	Herdsman	–
F103	DF9	–	Dairy farm	Herdsman	–
F104	DF9	–	Dairy farm	Crop manager	–
F105	DF9	–	Dairy farm	Herdsman, robot technician	–
F106	DF9	–	Dairy farm	In-house breeder	–
F107	DF9	–	Dairy farm	Hospital manager	–
F108	DF9	–	Dairy farm	Herdsman, equipment manager	–
F109	DF9	–	Dairy farm	Herdsman	–
F113	DF9	–	Dairy farm	Machinery operator	–
F143	DF11	Penny	Dairy farm	CFO / Farm co-owner	–
F144	DF11	Patricia	Dairy farm	CMO / Farm co-owner	–
F145	DF11	Priscilla	Dairy farm	COO / Farm co-owner	–
F146	DF11	Patrick	Dairy farm	Farm founder	50+
F147	DF11	Preston	Dairy farm	Farm manager	15+
F148	DF11	–	Dairy farm	Head herdsman	20+
F149	DF11	Pierce	Dairy farm	Herdsman	10+
F150	DF11	Pascal	Dairy farm	Finance officer	–
F151	DF11	–	Dairy farm	VP of Operations	–
F152	DF11	–	Dairy farm	Herdsman, robot technician	–
F153	DF10	Fidel	Dairy farm	Farm owner / manager	10+
F154	DF10	Fabricio	Dairy farm	Head herdsman, robot technician	–
F155	DF10	–	Dairy farm	Maintenance and equipment repair	–
F156	DF10	–	Dairy farm	Head herdsman	20+
F157	DF10	–	Dairy farm	Herdsman (hospital)	–
F158	DF10	–	Dairy farm	Herdsman	–
F159	DF12	Marco	Dairy farm	Farm owner / manager	10+
F160	DF12	Mike	Dairy farm	Robot barn manager	10+
F161	DF12	Markus	Dairy farm	Herd manager / hospital manager, in-house breeder	20+

ID	Site	Pseudonym	Organization	Role	Tenure (yrs)
F162	DF12	–	Dairy farm	Assistant manager	–
F163	DF12	–	Dairy farm	Milker	–
F165	DF13	–	Dairy farm	Farm owner	20+
F166	DF14	–	Dairy farm	Farm owner / manager	10+
F167	DF15	–	Dairy farm	Farm owner	40+
F168	DF16	–	Dairy farm	Farm owner / manager	10+
F169	DF17	Kevin	Dairy farm	Farm owner / manager	10+
F170	DF18	–	Dairy farm	Farm co-owner / manager	30+
F171	DF18	–	Dairy farm	Farm co-owner, herd manager	10+
F172	DF18	–	Dairy farm	Farm co-owner, finance manager	20+
F173	DF19	–	Dairy farm	Herd manager	20+
F174	DF19	–	Dairy farm	Lead herdsman	22
F175	DF19	–	Dairy farm	Farm owner	10+
F176	DF20	–	Dairy farm	Herdsman	25+
F177	DF20	Alex	Dairy farm	Farm manager	10+
F178	DF20	–	Dairy farm	Herdsman	–
F179	DF21	–	Dairy farm	Farm owner / manager	5+
F180	DF21	–	Dairy farm	Herd manager	5+
V1	VP1	Dr. E	Large animal veterinary practice	Veterinarian, partner	25
V2	VP1	Dr. S	Large animal veterinary practice	Veterinarian, partner	18
V3	VP1	Dr. B	Large animal veterinary practice	Veterinarian, associate	3
V4	VP1	Dr. W	Large animal veterinary practice	Veterinarian, associate	6
V5	VP1	Dr. C	Large animal veterinary practice	Veterinarian, partner	25
V6	VP2	Dr. K	Large animal veterinary practice	Veterinarian, partner	24
V7	VP3	Dr. M	Large animal	Veterinarian, partner	21

ID	Site	Pseudonym	Organization	Role	Tenure (yrs)
			veterinary practice		
V8	VP4	Dr. G	Large animal veterinary practice	Veterinarian, owner	22
V9	VP4	Dr. L	Large animal veterinary practice	Veterinarian, associate	3
V10	VP4	Dr. F	Large animal veterinary practice	Veterinarian, associate	2
V11	VP4	Dr. Y	Large animal veterinary practice	Veterinarian, associate	5
V12	VP5	Dr. D	Large animal veterinary practice	Veterinarian, owner	9

Notes. Informant IDs prefixed “F” denote farm-based informants; IDs prefixed “V” denote veterinarians. Site codes correspond to field sites listed in APPENDIX B. Tenure refers to years in the dairy industry as reported by the informant during fieldwork. All identifiers and organization names are pseudonymous.

APPENDIX D: Images from the Field

Images of Work



Image 1. Shadowing herdsmen during morning fetching to recruit cows who have been away from the robots for too long for milking.

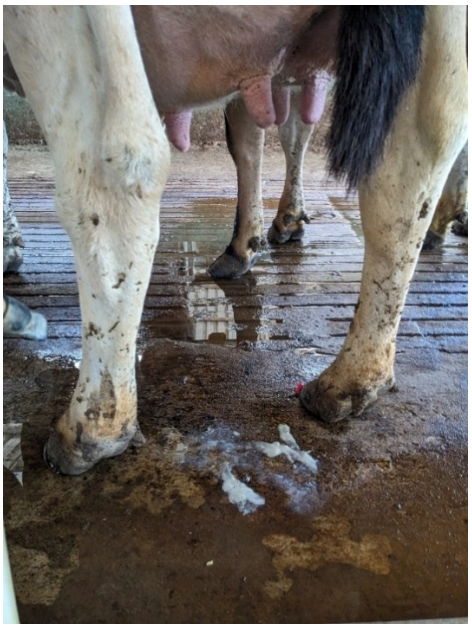


Image 2. During morning routine checks, herdsmen would often check for signs of mastitis, like chunky or discolored milk or discharge from teats. This image shows mucilaginous milk, a sign of infection.



Image 3. Milkers work around the clock to facilitate and monitor milk extraction. They also monitor for signs of mastitis, lameness, or other problems, often acting as the first line of defense to detect health issues.



Image 4. A safety sign posted in a milking parlor promoting social distancing during COVID-19. Considered essential services in food production, dairy farms continued operating throughout the pandemic.



Image 5. A breeder walks the cows pens daily and looks for evidence of being mounted. When cows are in heat, they allow other cows to mount them which rubs off the paint marks. When he finds a cow with rubbed off paint, he checks for other symptoms of heat, like clear discharge from the vagina, and artificially inseminates them if he judges them to be within the appropriate time window to conceive. He reapplies the paint every day.



Image 6. Sunrise at a farm during a routine veterinary herd check. Cows are restrained in head stanchions to hold them in position safely for examination.

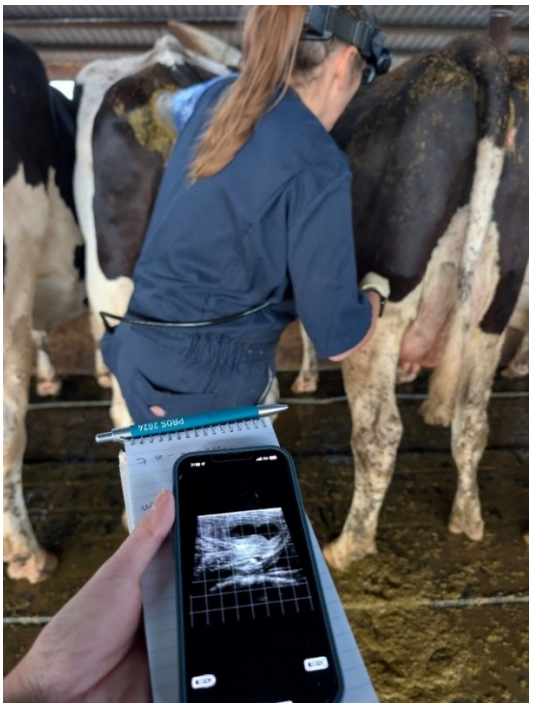


Image 7. A vet performs a rectal ultrasound, which sends an image to a head-mounted display and can also connect to a phone app via Bluetooth. Vets conduct rectal ultrasounds weekly to visualize the size of the fetus and diagnose conditions, such as fetal death, twins, or ovarian cysts.

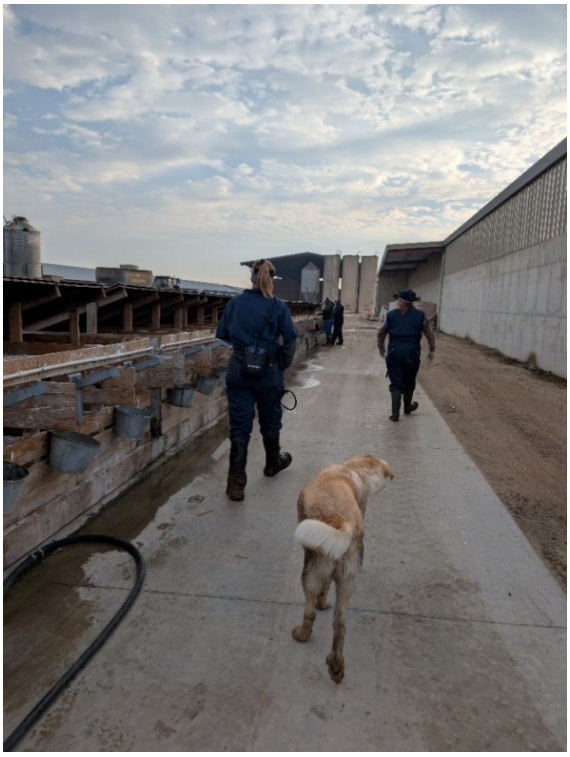


Image 8. Vets conduct a routine check of calf hutches accompanied by the farm's resident dog.

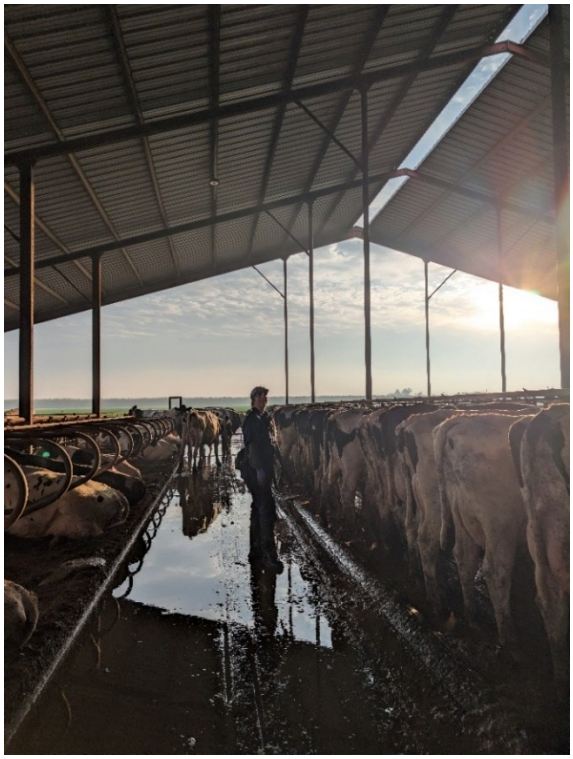


Image 9. A vet waits for instruction from a herdsman.



Image 10. A vet, herdman, and breeder consult about a cow. Herdsmen, veterinarians, managers, breeders, and other experts often work together during herd checks, coordinating activities like hormone injections and vaccines, exchanging information, examining cows, asking questions, and consulting one another when making decisions.



Image 11. A view from the cabin of a feed tractor as it dispenses a special feed forumulation designed by a nutritionist. Feed make up around 40% of the overall costs on a farm.



Image 12. An elevated view of a typical layout for a robotic barn. The waiting area is separated from the resting and feeding area, requiring cows to move through the robotic milkers to return to the pen.



Image 13. A vet walks a row of calf hutches at a calf ranch to monitor for signs of sickness.



Image 14. Older calves are housed together in aggregate pens. Vets walk these pens to monitor for signs of illness outbreaks which can spread quickly when calves are first introduced to free pens.



Image 15. Workers collect milk samples from cows that are suspected of having mastitis and prepare milk culture plates to incubate overnight. Positive bacteria growth is recorded in herd management systems (e.g., DairyComp) and the cows are separated from the milk herd and transferred to the hospital pen to be treated with antibiotics.



Image 16. A worker performs routine daily maintenance and calibration on a Lely A5 Astronaut milking robot.

Images of Milking Systems



Image 17. Example of a conventional parallel milking parlor. Cows are herded to a central building where milkers attach milking machines to each individual cow. Milk volume is measured by aggregating milk production for up to 250 cows per pen.



Image 18. The internal hardware and chemical analyzers in a Lely A5 Astronaut robotic milker. These systems are equipped with milk analyzers that measure and record dozens of metrics per udder quarter per milking.



Image 19. A Jersey cow peaks under the milking stall of a robotic milker to investigate the stranger while she is being milked.



Image 20. A DeLaval VMS (voluntary milking system) milking a Holstein.



Image 21. Cows are allowed to voluntarily visit milking robots based on milking permissions set by farm management.

Images of Data Systems and Platforms

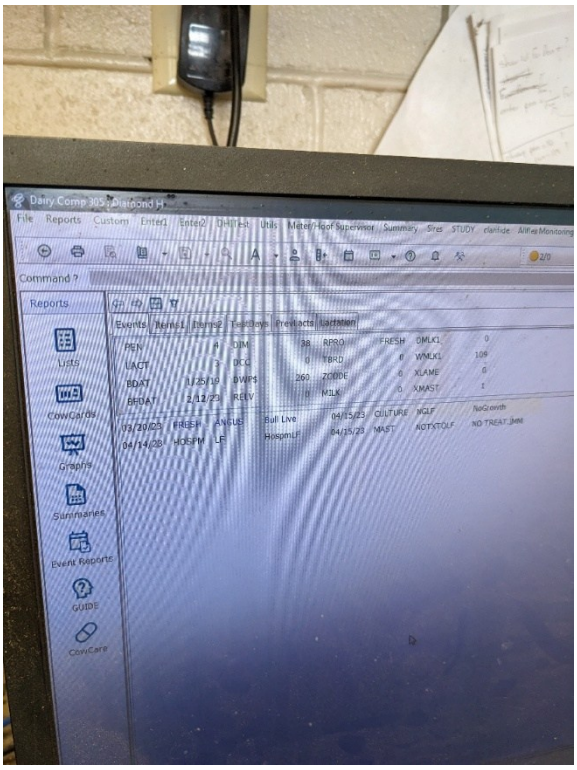


Image 22. An example of a cow's digital record in DairyComp.

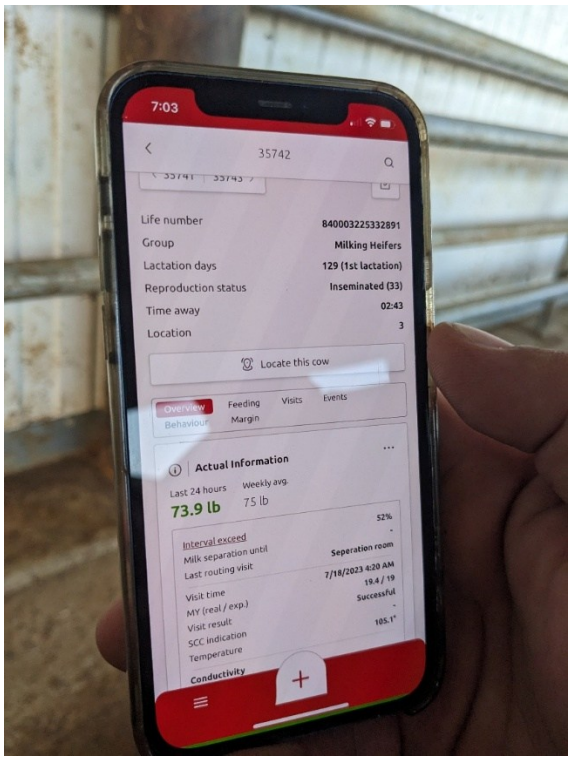


Image 23. An example of a cow's digital record in the Horizon phone app.

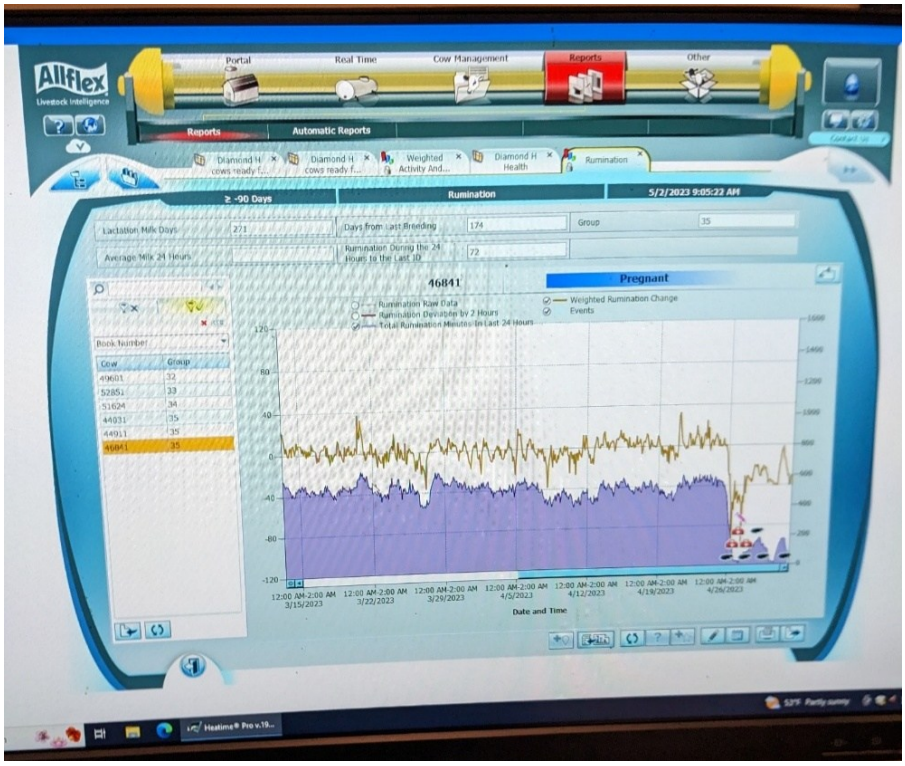


Image 24. An example of an activity graph displayed in AllFlex, a software program that encodes and analyzes data collected through activity tracking collars.



Image 25. An example of an activity graph displayed in Lely's Horizon program that encodes and analyzes data collected through activity tracking collars

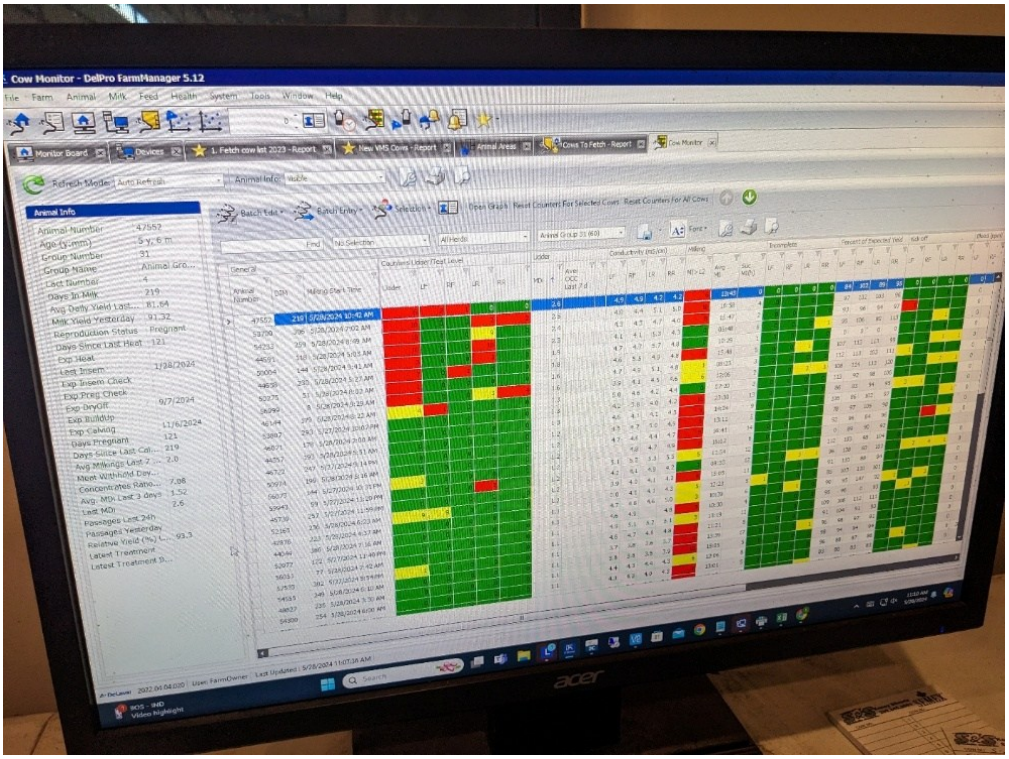


Image 26. An example of a mastitis report in DeLaval's DelPro FarmManager program showing somatic cell counts, conductivity, milk yields, kickoffs, and other metrics related to milk quality and milking efficiency. Cells marked in green are within normal parameters, yellow indicates borderline concern, and red denotes irregular parameters outside expected bounds.

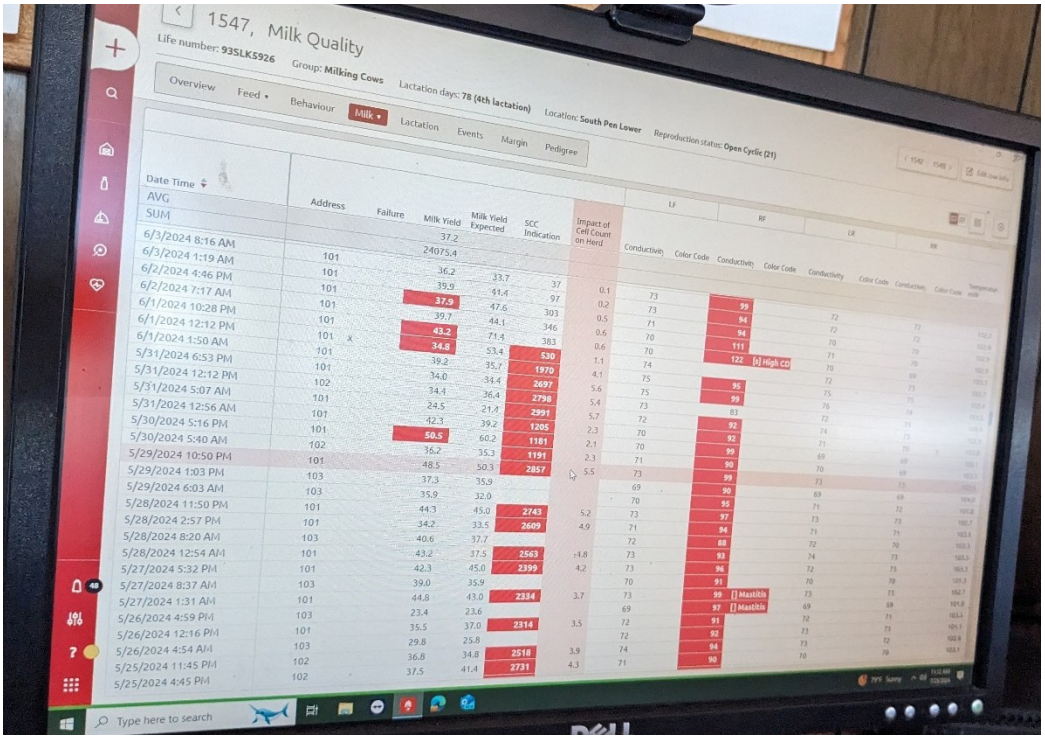


Image 27. An example of a milk quality report from Lely's Horizon program that shows somatic cell count, conductivity, milk yields, and other metrics related to milk quality and milking efficiency. Cells marked in red denote irregular parameters outside expected bounds.

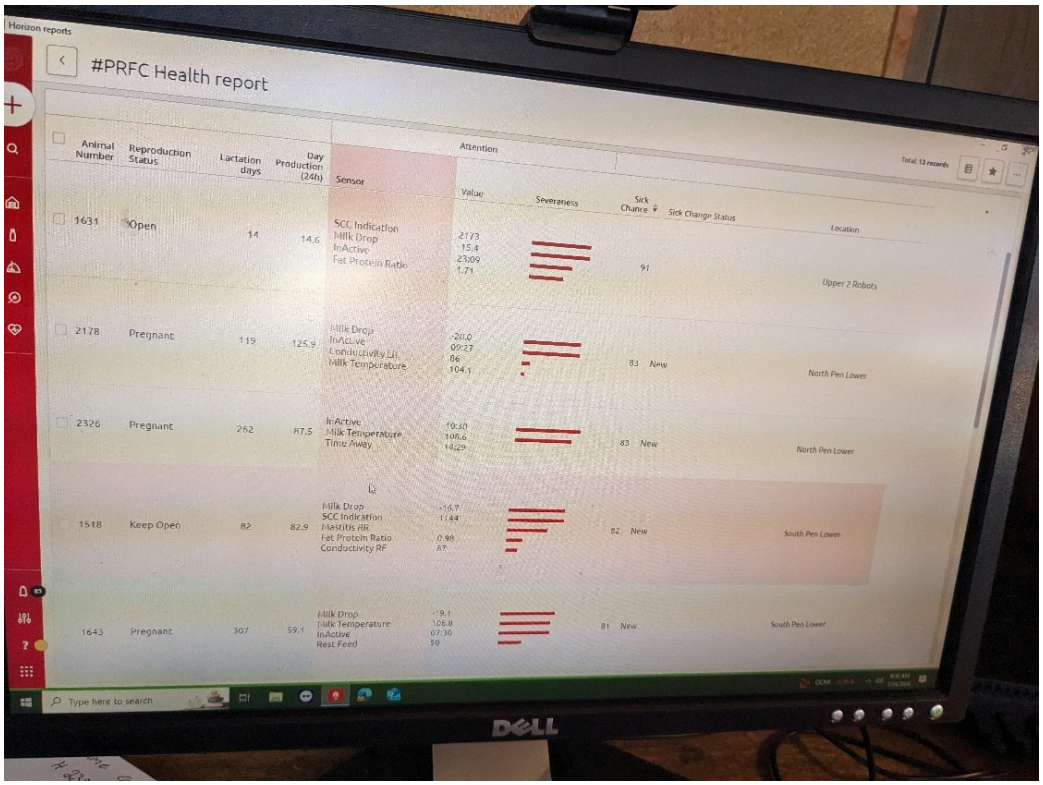


Image 28. An example of a health report generated in Lely's Horizon program that shows a list of cows with predicted illness. The report includes a sick chance score for each cow and the relevant measurements the algorithm used to predict this probability.



Image 29. Visual dashboard at a farm showing real-time performance data about all milking robots on the farm.



Image 30. A touchscreen display on the face of a DeLaval VMS robot that shows milk flow and yields per milking cup yielding real-time data for each udder quarter.

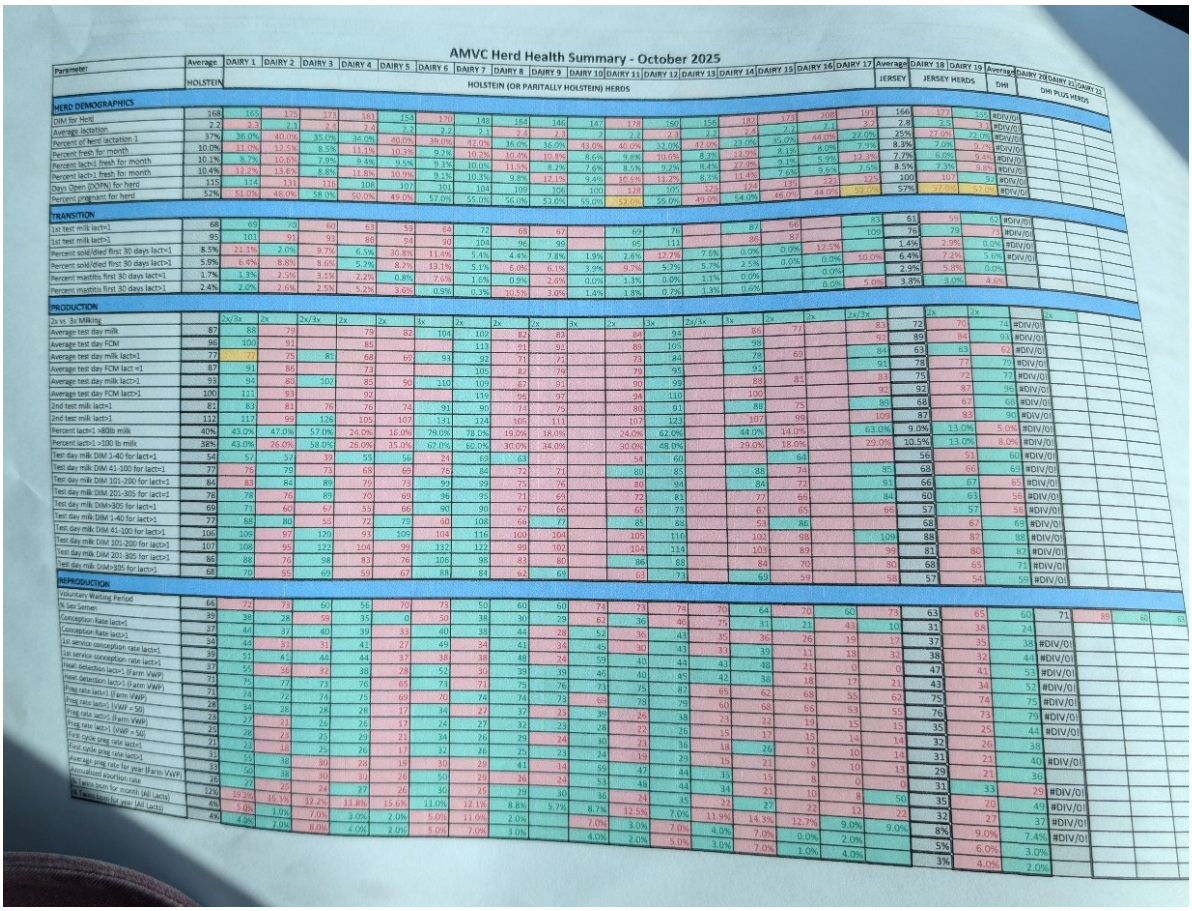


Image 31. A benchmarking spreadsheet compiled by veterinarians that gives each of their client an anonymized overview of their key performance indicators and how they compare with other client farms.

Images of Cows and the People Who Love them



Image 32. Jersey cows investigating an unfamiliar human on the farm.

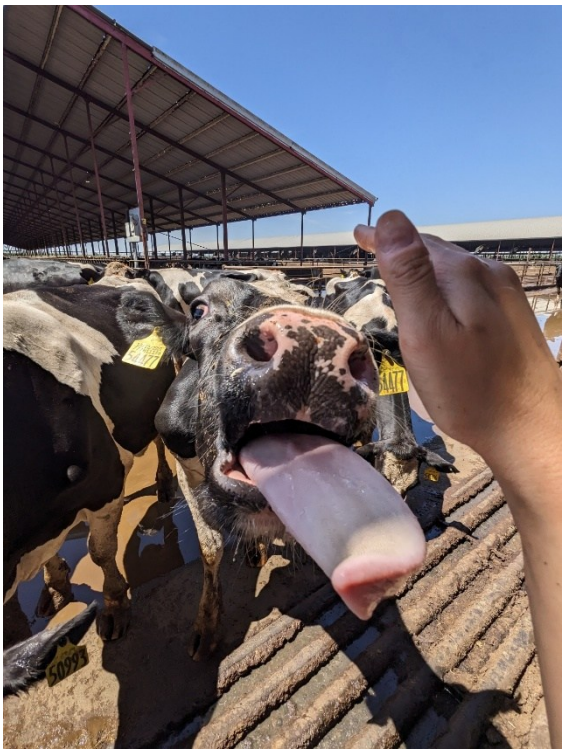


Image 33. Photographic evidence that cows will, indeed, “lick you to death.”



Image 34. Two researchers take a break from collecting data using a relational ethnographic approach.



Image 35. A selfie taken while shadowing herdsmen while they move cows to the milking parlor.



Image 36. A Holstein cow enjoying a back scratch from an automatic brush.



Image 37. An example of a cow with physical signs of mastitis (i.e., red, swollen right hind udder quarter).



Image 38. A cow in the maternity pen shortly after giving birth and her hour-old calf.



Image 39. Two researchers taking a break to document field work with a group of curious Jersey cows looking on in the background.



Image 40. A young calf in a pen being admired by a cow enthusiast.

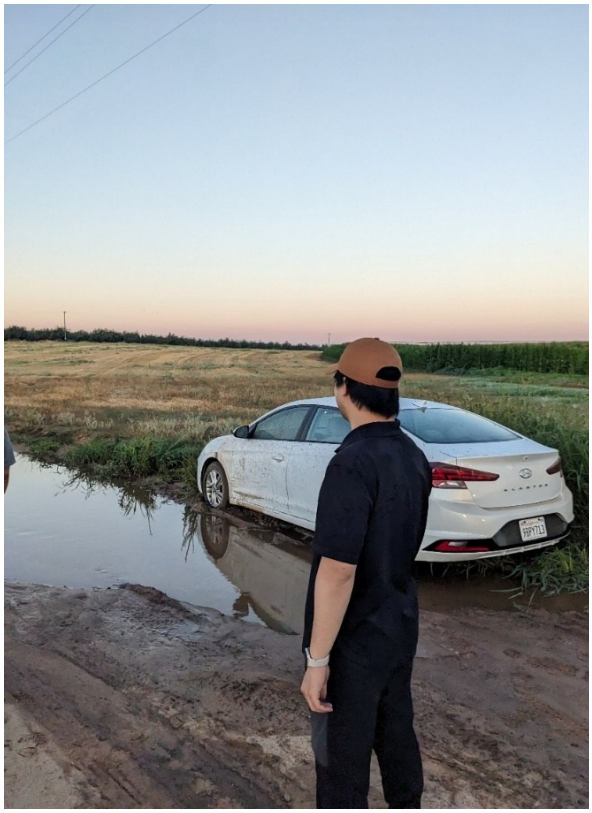


Image 41. A less-than-ideal scenario where a researcher's car got stuck in the mud driving onto a farm. Herdsmen arrived shortly to pull him out.



Image 42. A researcher gaining some “hands-on” experience with a Holstein calf.



Image 43. A researcher being investigated while shadowing routine work.



Image 44. A Holstein cow in the waiting area in line to be milked.



Image 45. A week-old brown Holstein calf in a calf hutch.