Check for updates

CASE STUDY

REVISED Trade-off hacking: Have the cake and eat it too? Turning competing interests into win-wins in data-driven food technology

[version 2; peer review: 1 approved, 2 approved with reservations]

Ferdinand Ferroli

Identity Valley Research, Unkel, 53572, Germany

v2 **First published:** 24 Nov 2025, **5**:358
https://doi.org/10.12688/openreseurope.21771.1
Latest published: 24 Feb 2026, **5**:358
https://doi.org/10.12688/openreseurope.21771.2

Abstract

Data-driven innovations in the food sector, from personalised nutrition to supply chain tracking, promise clear benefits but introduce complex trade-offs between competing interests such as personalisation and privacy, or performance and explainability. These tensions can hinder responsible innovation if addressed as zero-sum conflicts. This paper introduces the concept of "Trade-off Hacking", a user-centric technology design approach that reframes competing interests as opportunities for innovation striving for win-win outcomes. The concept is used as an analytical device to examine the practices of eight pilot projects funded by the Horizon Europe project "DRG4FOOD". This study applies an exploratory multiple-case study design to eight DRG4FOOD pilots selected through two Horizon Europe open calls (164 admissible proposals; 8 funded). Data sources included pilot applications, implementation documents, and project-level monitoring material generated during the incubation period. Analysis followed two steps: cross-case identification of recurring trade-offs; and pattern identification of resolution strategies into two emerging categories. Trough analysis of these real-world food tech pilots, this study moves beyond merely acknowledging trade-offs to identifying reproducible design patterns that rebalance, or even resolve them. The analysis focuses on three trade-offs (privacy/personalisation, performance/explainability, security/user experience) because these were the most recurrent and comparable across pilots; other tensions are discussed as out-of-scope for the primary framework. The analysis reveals a spectrum of strategies to achieve a rebalancing or resolving, from governance-based user controls and privacy-preserving architectures to co-design

Open Peer Review

Approval Status ? ✓ ?

	1	2	3
version 2 (revision) 24 Feb 2026		✓ view	? view
version 1 24 Nov 2025	? view	? view	

1. **Charis M Galanakis**, Galanakis Laboratories, Chania, Greece
2. **Juan Manuel Vargas-Canales**, Universidad de Guanajuato, Salvatierra, Mexico
3. **Agnieszka Kurdyś-Kujawska**, Koszalin University of Technology, Kwiatkowskiego, Poland

Any reports and responses or comments on the article can be found at the end of the article.

methodologies. The paper groups these solutions into two main categories: technology-driven resolutions, which use e.g. architectural or cryptographic methods to influence a trade-off, and cooperation-driven resolutions, which reframe value tensions as socio-technical negotiations.

Keywords
digital ethics, responsible innovation, privacy, cybersecurity, software development, personalisation



This article is included in the Horizon Europe gateway.

Corresponding author: Ferdinand Ferroli (f.ferroli@identityvalley.org)

Author roles: Ferroli F: Conceptualization, Writing – Original Draft Preparation, Writing – Review & Editing

Competing interests: No competing interests were disclosed.

Grant information: This project has received funding from the [European Union's Horizon 2020 research and innovation programme] under grant agreement No [101086523]
The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Copyright: © 2026 Ferroli F. This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

How to cite this article: Ferroli F. **Trade-off hacking: Have the cake and eat it too? Turning competing interests into win-wins in data-driven food technology [version 2; peer review: 1 approved, 2 approved with reservations]** Open Research Europe 2026, **5**:358 https://doi.org/10.12688/openreseurope.21771.2

First published: 24 Nov 2025, **5**:358 https://doi.org/10.12688/openreseurope.21771.1

REVISED **Amendments from Version 1**

In this revised version, the method is made clearer by stating more directly that the study is an exploratory multiple-case analysis and by explaining case selection, data sources, and analysis steps in a more straightforward way. The discussion of trade-off hacking is expanded, the focus on three trade-offs is justified more clearly, and a more critical section is added on limitations, replication, and context. The conclusions are also toned down and now stress that trade-offs are often reconfigured rather than fully resolved.

Any further responses from the reviewers can be found at the end of the article

Introduction

Digital technology is increasingly reshaping not only the global food system but also the relationship towards the food sector. Innovations ranging from AI-driven personalised nutrition applications to blockchain-based traceability systems offer the potential for promoting public health, empowering consumers, and promoting sustainability (see Ellahi *et al.*, 2023; Wu *et al.*, 2025). However, the integration of data-driven technologies into this critical sector also introduces important ethical and practical conflicts. An application that personalises meal plans based on sensitive health data must simultaneously protect user privacy; an algorithm that predicts the carbon footprint of food products with great precision may not be trusted by retailers and regulators if its calculations are not transparent. These tensions are not merely technical hurdles but represent value conflicts at the core of responsible innovation (Detopoulou *et al.*, 2023; Donovan *et al.*, 2025; Reis *et al.*, 2021).

Addressing these challenges is a central mission of the Horizon Europe "Digital Responsibility Goals for Food" (DRG4FOOD) project. With a total funding of €4 million, DRG4FOOD aims to foster a data-driven food system that inspires trust throughout the digitalised food chain. The project functions as a responsible technology incubator, having distributed €1.9 million through two Open Calls to eight pilot consortia developing data-driven food tech applications (DRG4FOOD, n.d.). Condition for pilot consortia applying to the incubator programme was the implementation of/adherence to the project's guiding framework of the seven Digital Responsibility Goals (DRGs). These principles – including Digital Literacy, Cybersecurity, Privacy, Data Fairness, Trustworthy Algorithms, Transparency and Human Agency & Identity – provided the normative backbone for the pilot projects (Identity Valley, 2025; Meier *et al.*, 2022). This governance structure enabled a real-world laboratory where diverse teams were explicitly tasked with navigating the inherent tensions of building responsible technology while competing under real-world market pressures.

The DRG4FOOD programme design and coaching likely influenced project choices, however, this study does not infer direct causality between incubator support and observed design outcomes. By establishing digital responsibility as a key selection criterion for pilots and providing ongoing coaching and resources, the project aimed to create a distinct ecosystem conducive to responsible innovation (Ferroli *et al.*, 2023). This environment incentivised and enabled teams to "invest" in resolving complex trade-offs, rather than defaulting to the path of least resistance often dictated by requirements of the market alone. The advanced solutions documented in the pilots – from cryptographic protocols of the consortium "PINACLE" to the comprehensive open-source governance of "Nutrisight" – suggest that the programme's structure catalysed engagement with responsible design principles (Ferroli & Weich, 2025). In this context, this "incubator effect" helps understand the successful pilots not as isolated instances, but as potential outcomes of a deliberately created environment for responsible innovation.

It follows that by analysing the practical strategies of the DRG4FOOD pilots through the conceptual lens of "trade-off hacking" we can identify replicable patterns that can help resolve competing interests in other contexts. The evidence from these eight pilot cases suggests that ethical and functional requirements can be designed to be mutually reinforcing rather than oppositional. While achieving a perfect win-win for every stakeholder in every scenario remains a challenging ideal, this analysis moves beyond merely acknowledging trade-offs to conceptualising a constructive empirically grounded path forward for the development of trustworthy and innovative technology for the food sector and beyond.

This paper treats DRG4FOOD as a bounded empirical setting in which teams operated under shared responsibility criteria, not as proof that one funding model universally causes better outcomes. The objective is to analyse how teams navigated value tensions under these conditions and to assess which patterns may transfer to other settings.

Background

To systematically analyse the paths taken by the DRG4FOOD pilots, this paper proposes the concept of "trade-off hacking". This section defines the term, situates it within academic discourse, and outlines the analytical framework used in the subsequent analysis.

Defining "trade-off hacking"

Trade-off hacking is a user-centric approach to digital technology development that reframes a seeming either-or dilemma into an opportunity for innovation. It challenges the assumption that certain design goals are mutually exclusive. Rather than accepting a compromise that weakens both sides of a dilemma – for instance, moderately good personalisation with moderately weak privacy – trade-off hacking seeks to reconfigure the problem space to satisfy both competing interests as much as possible. This approach transforms a zero-sum conflict into a positive-sum innovation challenge. While a perfect positive-sum outcome for all stakeholders (e.g. users, developers, and businesses) may be not achievable, this approach focuses on reconfiguring the problem to move beyond simple compromise and facilitate new, more optimal solutions.

Hence, in this paper, trade-off hacking does not imply that all conflicts are fully resolvable. Some tensions remain structurally persistent, and improvements on one dimension can introduce costs in latency, usability, implementation burden, or governance. A successful "hack" is one that improves outcomes for the most relevant stakeholders in a specific context, without claiming to always benefitting everyone equally.

Situating the concept in academic discourse

The concept of trade-off hacking is inspired by and builds upon several established theoretical frameworks in technology design and digital ethics.

The concept is an attempt at a practical application of Value-Sensitive-Design (VSD), a methodology that advocates for the proactive and principled integration of human values into the entire technology design process (Friedman *et al.*, 2008). The dilemmas faced by the DRG4FOOD pilots – as with many product designers and software developers – are conflicts between competing values, such as the value of privacy versus the value of convenience, or transparency versus security. VSD provides the more theoretical frame to address these values explicitly, and trade-off hacking offers a more pragmatic strategy.

Furthermore, the principles of trade-off hacking are also reflected in the emerging field of Value-based Engineering (VBE). Building on the foundational work of VSD, VBE offers a more structured and formalised methodology for integrating ethical values into system design, grounded in the IEEE 7000 standard. Unlike VSD, which provides a broader theoretical framework, VBE focuses on a traceable, step-by-step process that translates identified stakeholder values into concrete technical and organisational requirements, known as Ethical Value Requirements (EVRs) (Spiekermann & Winkler, 2022). By drawing from VSD, co-design, and software engineering, VBE provides a practical pathway for ensuring that ethical considerations are not just discussed but are systematically embedded and verified throughout the development life cycle.

Aspects of trade-off hacking are also intertwined with participatory and co-design methodologies. These approaches involve stakeholders, particularly end-users, as active partners in the design process (Sanders & Stappers, 2008). As evidenced by several DRG4FOOD pilots, engaging users and domain experts directly is a powerful strategy for identifying and resolving value tensions. Stakeholders become allies in innovation, offering insights that can reveal novel, win-win solutions that a purely technical/product team might overlook. This transforms the resolution of trade-offs from a top-down technical decision into a collaborative, socio-technical negotiation (Sadek & Mougenot, 2024; Steinke *et al.*, 2021).

At the foundation of the DRG4FOOD approach stand the Digital Responsibility Goals (DRGs), a framework that reframes society's relationship with technology by linking digital innovation to human values and trust (Meier *et al.*, 2022). Much like the United Nations' Sustainable Development Goals (SDGs) united global actors around a shared sustainability agenda, the DRGs provide a common orientation for a trustworthy digital transformation. Comprising seven interrelated goals – Digital Literacy, Cybersecurity, Privacy, Data Fairness, Trustworthy Algorithms, Transparency, and Human Agency & Identity – they serve as measurable reference points for responsible digital behaviour. In DRG4FOOD, this framework functioned as both governance scaffold and design compass, enabling project teams to translate ethical imperatives into concrete design requirements. The DRGs thus do not merely define abstract ideals but offer a normative structure through which tensions such as privacy versus personalisation or performance versus explainability can be more systematically addressed.

Establishing the analytical framework of trade-offs

The analysis in this paper is structured around three key trade-offs that are recurrent in the literature and were prominently featured in the experiences of the DRG4FOOD pilots. These three dilemmas serve as thematic pillars for the multi-case analysis:

Privacy vs. Personalisation: The tension between the need for user data to deliver tailored services and the imperative to protect users' sensitive personal information (see e.g. Awad & Krishnan, 2006; Cloarec, 2020).

Accuracy/Explainability vs. Performance: The conflict where the most accurate and powerful algorithmic/AI models, such as those based on deep neural networks, are often the most opaque and difficult to interpret, creating barriers to trust and accountability (Crook *et al.*, 2023; Van Der Veer *et al.*, 2021).

Security vs. User Experience: The friction where essential security and compliance measures, such as complex authentication or lengthy consent forms, can create a cumbersome and frustrating user experience, even leading users to circumvent them (Jacobs & McDaniel, 2022).

These three key trade-offs were selected using three criteria: recurrence across pilots, direct relevance to DRG-aligned design decisions, and sufficient evidence for cross-case comparison. Other tensions (e.g., cost vs. inclusivity, or speed-to-market vs. participatory design) were observed but occurred less consistently or lacked enough comparable evidence across all pilots. This framing improves analytical focus, however, also narrows the generalisability of conclusions.

Methodological approach

This study adopts an Exploratory Case Study approach (cf. Mills et al., 2010, 372f). An exploratory design is appropriate because "trade-off hacking" is still conceptually emergent and the goal is analytic clarification and pattern identification rather than hypothesis testing. Cases were included if they developed a data-driven food technology, engaged at least one DRG-relevant value tension, and provided sufficient documentary evidence for cross-case analysis. Data sources included pilot applications, implementation documents, and project-level monitoring material generated during the incubation period. For each case, trade-offs, implemented mitigation strategy and enabling conditions were identified. Then patterns were matched across cases to classify strategies as technology-driven or governance-driven.

Trade-off hacking: vignettes from the DRG4FOOD pilots

The eight pilot projects selected and supported by the DRG4FOOD consortium provide for some valuable "micro" case studies to observe how trade-offs are navigated in practice. The DRG4FOOD open call selection process for the pilot projects asked applicants to specify in their proposal not only the market viability and scalability of their solutions but decidedly also their measures to achieve digital responsibility (as defined by the DRGs, see Identity Valley, 2025; Meier *et al.*, 2022). In total 164 proposals between two calls were admissible with an average score of 3,6 (call #1) and 3,7 (call #2) out of five for criteria of digital responsibility (Ferroli & Weich, 2025). This additional requirement for digital responsibility was well-received and understood by applicants, nudging them into a situation where they needed to confront design conflicts already from ideation and navigate trade-offs they otherwise might not have.

The following micro case studies, or vignettes, are not presented as fundamental innovations in responsible technology but rather as small-scale, practical examples to underpin the argument that the deliberate engagement with design tensions yields replicable strategies to resolve them. The pilots exhibit progress in rebalancing these trade-offs, even if they cannot completely resolve them. A comparative analysis of strategies organised by the key trade-offs identified will be presented in this section. Table 1 provides a high-level overview of the pilots and the primary trade-offs they faced and addressed.

Table 1. Overview of DRG4FOOD Pilot Projects and Key Trade-Offs.

Project Name	Core Functionality	Key Trade-Off
ATTESTED	Farm-to-fork traceability system for small producers using IoT, RFID, and portable sensors.	Security vs. User experience
Cacao-Tech	Cacao quality and traceability platform using near-infrared technology and a tracking technology.	Privacy vs. Personalisation (for framers); Security vs User Experience
DISH	Personalised recipe application using machine learning and symbolic AI, focusing on user privacy.	Privacy vs. Personalisation; Explainability vs. Performance
GENIE	Ultra-personalised nutritional recommender based on genetic, gut microbiota, and blood test data.	Privacy vs. Personalisation
NutriSight	AI tool to automatically extract nutritional information from food packaging photos for Open Food Facts.	Explainability vs. Performance

Table 1. *Continued*

Project Name	Core Functionality	Key Trade-Off
NutriWell	AI-based personalised nutrition platform for elderly individuals, with a focus on social inclusion.	Privacy vs. Personalisation; Explainability vs. Performance
PINACLE	AI-driven nutrition recommender matching donated food with recipients' dietary needs using blockchain.	Privacy vs. Personalisation
SafeNutriKids	AI-driven personalised nutrition education app for children aged 6–12.	Security vs. User Experience; Privacy vs. Personalisation

Reconciling personalisation and privacy: from data collection to user sovereignty

The "personalisation-privacy-paradox" describes the tension where services need (sensitive) personal data to be effective, yet users are wary of sharing it due to privacy concerns (Chellappa & Sin, 2005; Xu *et al.*, 2011). The DRG4FOOD pilots showcase a spectrum of strategies to resolve this dilemma, moving beyond simple consent models toward true user sovereignty. These solutions can be categorised along a continuum of intervention, from governance-based controls to fundamental architectural and cryptographic redesigns.

At the governance and user interface level, the GENIE project provides a suitable example. Tasked with handling exceptionally sensitive genetic data the team implemented a multi-layered approach. The core strategy was to combine standard backend measures, such as anonymising user data and separating personal information from the analytical dataset, with a user-facing "data control panel". This panel empowers users to actively manage their preferences in detail and opt out of data sharing at any time, shifting the dynamic from passive consent to active control.

Moving to the architectural level, the DISH project demonstrates a different approach. To deliver personalised recipes, the team made deliberate architectural choice to store user profiles and persona information "on-device" only. By redesigning the data flow to avoid centralised collection of personal data, they effectively dissolve the privacy risk at its source. Personalisation is achieved through client-side processing, demonstrating that it is possible to deliver a tailored experience without extracting and aggregating sensitive data on a central server. This approach makes the trade-off less about balancing risks and more about eliminating them through system design.

A more technologically focused approach is found in the PINACLE project, which leverages cryptographic methods to make the trade-off obsolete. PINACLE, which matches donated food with recipients' needs, uses a combination of Self-Sovereign-Identity (SSI) principles, Verifiable Credentials (VCs), and Zero-Knowledge-Proofs (ZKPs). Zero-Knowledge-Proofs are a cryptographic protocol that allows one party to prove to another that a statement is true without revealing any information beyond the validity of the statement itself (Garimella & Conway, 2024). In PINACLE's context, this allows a food recipient to prove eligibility without disclosing sensitive information. This represents a significant shift: the conflict between verification (needed for personalisation/matching) and privacy is not balanced but is rendered virtually irrelevant. The system can verify a user's dietary needs or status without ever "seeing" the underlying sensitive data. Paired with VCs, which give users granular, portable control over their own data, this approach enables true user sovereignty (Sedlmeir *et al.*, 2021).

Together, these three cases illustrate that resolving the privacy dilemma is not a monolithic strategy. It exists on a spectrum: GENIE's approach is about governing data flows, DISH's is about re-structuring them, and PINACLE's is about cryptographially transforming the nature of verification itself. This provides a valuable framework for practitioners, who can select the level of intervention best suited to their context, resources, and the sensitivity of the data they handle.

Yet, while several pilots improved privacy protections, an open question is whether stronger safeguards may, in some contexts, reduce data availability for ongoing model improvement. Similarly, privacy-enhancing approaches (e.g., advanced cryptographic methods) may introduce integration demands not yet observable. This suggests the trade-off may be, in many cases, reconfigured rather than removed, with possible costs shifting from privacy risk toward technical and organisational effort.

Explainability and performance: Building trust in a black box

A central challenge in modern AI is that the most powerful predictive/generative modesl, particularly in deep learning, often function as "black boxes", making their "reasoning" opaque (Hassija *et al.*, 2023). This lack of transparency can be a major barrier to user trust and adoption, especially in important domains like health and nutrition. The pilots developed several strategies to provide explainability – and for that matter ensure accuracy – without impinging on performance, by, for example, resorting to less efficient algorithms.

The NutriWell project, which generates personalised nutrition plans for the elderly, demonstrates a proactive, user-centric approach to explainability. Their platform features "Info Points" strategically placed throughout the interface. These provide detailed, contextual explanations of the AI's recommendations, including insights into their sources, and the scientific foundations behind them, along with the direct links to relevant scientific publications. Furthermore, the platform visualises the precision of its AI by showing users the percentage deviation between the generated meal plan and their initial nutritional requirements. This approach builds trust not by simplifying the underlying model, but by making its outputs transparent, verifiable, and comprehensible to the user.

A different strategy is exhibited by NutriSight, which uses a neural model to extract nutritional data from images. The team resolved the explainability dilemma by implementing a robust "human-in-the-loop" validation system. Their setup ensures that every prediction is validated by a contributor before it is integrated into the Open Food Facts database. This is a pragmatic solution: It allows the project to leverage a high-performance, complex model, while relying on human oversight as the guarantor of accuracy and trustworthiness. Here, trust is not established through an explanation of the algorithm's internal workings but through the human verification process itself.

Finally, the DISH project chose an approach based on technology selection. The team uses methods from the domain of symbolic AI that allows reliable tracing of the computation process back to the original information sources. Unlike deep learning models, symbolic AI systems operate on more explicit rules and logic, making their "reasoning" inherently transparent and traceable (Attoresi, 2025; Yu *et al.*, 2023). In this case, the trade-off is addressed at the foundational level by selecting an AI paradigm where explainability is a native property. This approach prioritises transparency, potentially trading a degree of predictive performance for comprehensibility.

A further issue to examine is whether explainability gains could create operational constraints under real-world workload conditions. For example, human-in-the-loop validation may improve trust but could become a bottleneck as deployments scale. Future evaluations should therefore test whether these approaches deliver stable resolutions over time, rather than one-off improvements.

Integrating security and user experience

Security and regulatory compliance are essential for any digital service, but their implementation often introduces friction that can negatively impact user experience, leading to frustration or non-compliance (Lennartsson *et al.*, 2021). The pilots show that by treating security not as a technical checklist but as a user-centred design challenge from the beginning, it is possible to create systems that are both safe and seamless.

The SafeNutriKids project faced a version of this dilemma, as it was designed for children, a vulnerable user group requiring high standards of data protection and parental consent. Their solution was to reframe these stringent compliance requirements as an opportunity for user-centric design. Instead of presenting parents with legalistic consent forms, the app communicates its data policy to children through age-appropriate, interactive explanations, using visual cues and simple language. This approach transforms a legal necessity into an engaging and educational feature. It enhances security by fostering genuine understanding and informed consent. While simultaneously improving the user experience for both children and parents.

The Cacao-Tech project provides another example: The team's traceability platform encountered early concerns from cocoa buyers about sharing competitively sensitive data regarding their farmers. A purely technical solution – e.g. simply stating that the data is encrypted – might not have been sufficient to build trust in a relationship-based industry. The team's solution was therefore more social, committing to a participatory approach of co-designing/co-developing the tools with the users. By actively involving farmers, buyers, and other stakeholders in the design of the data sharing and security mechanisms, they ensure the final system is not only technically secure but also accepted and usable within the real-world context of the cacao industry. This co-design process pre-emptively alleviates the security-UX trade-off by ensuring the security model is one that users themselves have helped create, validate, and therefore feel ownership over.

However, while co-designing with users can build trust and credibility, it can also take significant time and facilitation capacity, which may be difficult for teams with limited resources. This means that, in practice, teams may need to balance the depth of participation with what is realistically feasible in their operational context.

Synthesis: Two patterns of trade-off hacking

The insights from the DRG4FOOD pilots allow for a synthesis that groups their practices into two overarching types of patterns for the resolution of conflicts: technology-driven resolutions and governance-driven resolutions. This framework helps classify the strategies used to transform design conflicts towards win-win outcomes, or at least into

Table 2. Overview of Trade-Off Hacking Patterns in DRG4FOOD Pilots.

Project Name	Key Trade-off	Resolution	Pattern Category
DISH	Privacy vs. Personalisation/ Explainability vs. Performance	On-device data storage and processing/Use of inherently transparent Symbolic AI	Technology-driven
PINACLE	Privacy vs. Personalisation	Zero-Knowledge Proofs (ZKPs) for verification	Technology-driven
GENIE	Privacy vs. Personalisation	User-facing data control panel for granular consent	Governance-driven
NutriWell	Explainability vs. Performance	"Info Points" in UI to explain AI recommendations	Governance-driven
SafeNutriKids	Security vs. User Experience	Interactive, age-appropriate consent mechanisms	Governance-driven
NutriSight	Explainability vs. Performance	Human-in-the-loop validation of all AI outputs	Governance-driven
Cacao-Tech	Security vs. User Experience	Co-design of data sharing rules with stakeholders	Governance-driven

significantly more favourable trade-offs than traditional compromises. While a more detailed categorisation could have been made, it became clear that the success of many resolutions is tied to many more factors which are specific to the application/use case at hand. In the light of this, the synthesis below – summarised in Table 2 – can be seen as a rough charting of the territory rather than a guideline for exact reference.

Technology-driven hacks

This category of patterns involves embedding the solution to conflicting interests directly into the technical setup of the system. These approaches rethink the underlying architecture, algorithms, or cryptographic protocols to eliminate the trade-off at its source. For example, the DISH project tackled the personalisation-privacy paradox by making fundamental architectural choice to store and process all user data exclusively on the user's device. This design dissolves the privacy risk of data transfers and storage, demonstrating that a tailored user experience does not require a sacrifice of personal data.

Taking this principle a step further, the PINACLE project employed advanced cryptographic methods like Zero-Knowledge-Proofs (ZKPs). This technology allows the system to verify a user's eligibility for a food donation without ever needing to see the sensitive personal data that proves eligibility, rendering the conflict between verification and privacy risks obsolete. The choice of algorithm itself can also be a successful resolution. DISH again provides an example by opting for symbolic AI, whose rules-based nature makes its "reasoning" inherently transparent. This choice prioritises explainability, addressing the "black box" problem before it arises, rather than trying to explain an opaque model after the fact. In each case, the resolution is not a feature added on top of the system, but it *is* the system.

Governance-driven hacks

This category reframes value tensions as socio-technical challenges that are best solves through better rules, processes, user controls, and stakeholder collaboration. These solutions build trust not just through code, but though transparency, agency, and shared ownership of how the technology operates. A common pattern is empowering users with direct control. The GENIE project, handling extremely sensitive genetic and health data, implemented a user-facing "data control panel" that allows individuals to actively manage their sharing preferences, turning passive consent into active data governance. This focus on clarity and agency is also visible in how pilots approached the provision of complex information. Instead of dense legal text, SafeNutriKids transformed consent forms into age-appropriate, interactive explanations, making compliance more of an engaging experience. Similarly, NutriWell used "Info Points" within its interface to offer clear, contextual justifications for its AI's nutritional advice, building user trust through transparency.

Beyond individual user interfaces, governance-driven resolutions can also shape the operational process. The NutriSight project, for instance, uses AI to extract data but govern its use with a strict "human-in-the-loop" protocol, where every automated entry is validated by a person. Here, trustworthiness comes from the process, not the algorithm. The collaborative principle is also present in the Cacao-Tech project, which resolved data-sharing fears among competitors

by co-designing the system with its stakeholders. By making the design process itself a "negotiation", the final technology is shaped with the trust of its stakeholders.

Limitations

Technology-driven patterns are strongest when risks can be technically specified and teams have advanced engineering capacity. However, they often require higher investment, specialised skills, and longer integration cycles. Governance-driven patterns can be faster to implement and improve legitimacy, but their effectiveness depends on participation quality, competence, and institutional trust. In practice, the most robust cases might be hybrid: governance to define acceptable constraints, followed by targeted technical enforcement.

Discussion

A central limitation is the "incubator effect": pilots operated under funding, coaching, and responsibility-oriented selection criteria that are not typical of many commercial environments. Without these supports, teams may prioritise speed-to-market and short-term revenue over resource-intensive responsible-design measures.

Replication outside Horizon Europe is plausible but conditional on enabling factors: dedicated budget for responsible design, increased governance capacity by facilitators of framework programmes for research and development, cross-disciplinary expertise, and stakeholder access. Applicability is likely lower in low-margin markets, low-trust institutional contexts, or organisations with limited technical maturity.

Yet, the analysis of the DRG4FOOD pilots and the emergent patterns of trade-off hacking carry some useful implications for practitioners, policymakers, and the financial institutions funding digital innovation.

The role of the responsible technology incubator

Above insights suggest that the structure of the DRG4FOOD project itself played a role in fostering these innovative solutions. By making digital responsibility a primary criterion for selection and funding, the program highlighted its importance and created a meaningful incentive. The provision of coaching and resources further provided the capacity for teams, many of whom are SMEs or startups, to invest time and their own resources into solving the value tensions present. This "responsible incubator effect" suggests that funding mechanisms like Horizon Europe can be powerful levers for promoting responsible practices and innovation: To cultivate trade-off hacking, such programs should not only fund technology development but also explicitly reward and support the interdisciplinary and user-centred processes required to resolve value tensions. This de-risks the investment in responsible innovation, which might otherwise be sidelined in favour of more immediate commercial goals.

Implications for practitioners

For technology developers, the central takeaway is that while trade-offs can rarely be completely resolved to achieve a win-win for all stakeholders, the pursuit of a better resolution is a powerful driver of innovation, and potentially adoption. The "Edelman Trust Barometer Special Report: Tech Sector" series of global surveys (n = 32.000+) have consistently found that trust is tied to adoption. In their 2024 report, for example, they found that the largest barrier to AI adoption are concerns about privacy (Edelman, 2024, p. 28). However, while a solution like on-device processing is a clear win for user privacy, it may not be a win for a business whose revenue or innovation strategy depends on collecting centralised data. This is exactly what "trade-off hacking" can be valuable as a way of thinking. It nudges practitioners to constantly seek better optima. In the above example, a next-level solution like federated learning might bring the desired win-win, as data stays localised on user devices while the business can still profit from aggregated model insights without infringing on individual privacy (Loftus *et al.*, 2022). Therefore, the goal is to cultivate this mindset, viewing trade-offs not as constraints but as "innovation prompts".

The patterns identified in this paper – from user-centric control panels to privacy-preserving cryptography – provide a more nuanced toolkit for responding to these prompts. This requires a dual skillset: the technical know-how to identify and implement technology-driven solutions and the "social skills" to guide governance-driven resolutions through co-design and transparent communication. Adopting an interdisciplinary approach, bringing in expertise from ethics, law, or social sciences in addition to engineering, is most effective for identifying and implementing those more holistic solutions.

Development teams should prioritise technology-driven strategies when data-related risks are high, the technical architecture can be adapted, and there is sufficient long-term capacity to maintain and update solutions. Governance-driven strategies might be more appropriate when values are contested, workflows are highly context-dependent, or rapid legitimacy-building with stakeholders is critical. Under real-world constraints, a hybrid approach might often be most

effective: begin with governance (e.g., problem framing, consent design, and accountability arrangements), then introduce targeted technical controls where risks remain. Across contexts, decision-making in these cases could usefully be informed at least by expected scale, implementation cost, organisational maturity, and regulatory exposure.

Implications for policy-makers

The experiences of the pilots indicate that regulations like the EU General Data Protection Regulation (GDPR), while sometimes perceived as burdensome, can act as powerful catalyst for innovation. The need to comply with stringent privacy and consent requirements pushed teams like SafeNutriKids and PINACLE to develop solutions that not only meet the legal requirements but also enhanced user trust and thereby created a USP and potential competitive advantage. This suggests that policy should not only set high standards for digital responsibility but also actively support the ecosystem needed to meet them. This includes fostering environments where innovators are encouraged and equipped to pursue the ongoing process of optimising trade-offs, funding research into privacy-preserving technologies, promoting the dissemination of successful design patterns, and fostering environments where innovators are encouraged and equipped to find smart, user-centric implementations of regulatory requirements.

Future research

This investigation is based on an analysis of eight pilot projects within a specific, EU-funded programme. While this provides a coherent dataset, the findings are not generalisable to all contexts and serve primarily as a mental stepping stone to facilitate deeper conversations about re-thinking digital technology development. Future research should include longitudinal studies to track the long-term success and scalability of these "hacked" solutions in the market. Comparative analyses of projects developed outside of a "responsible incubator" context is needed to determine the real extent to which these patterns are adopted organically. Further investigation into the practical implementation of the emergent patterns, particularly the organisational and technical requirements for implementing effective technology- or governance-driven "hacks", would be a rewarding path for future work. Also, there is a need to develop and apply harmonised outcome indicators across comparable digital solutions so that cross-case analysis can move beyond qualitative pattern mapping and assess relative effectiveness more robustly.

Conclusion

In the pursuit of a data-driven food system that is sustainable, healthy and fair, it is tempting to view digital technology as a rigid means to an end: that we must choose between personalisation and privacy, performance and transparency or good user experience and security. This paper has argued that this framing often stifles innovation needed to move forward in creating a more user-centric digital ecosystem – and therefore also realise the ideals of a functioning data-driven food system.

Adopting a mindset of "trade-off hacking" – an intentional, creative and human-centred approach to digital technology development – opens new pathways to convert these either-or choices into both-and opportunities.

The eight pilot cases analysed indicate that trade-offs can often be reconfigured, but rarely entirely resolved, and usually at the cost of additional technical, organisational, or governance complexity. Trade-off hacking is therefore best treated as a "constrained design heuristic" rather than a universal recipe for win-win outcomes. Future studies should test these patterns in non-incubator settings and report failed repliactions as systematically as successful adaptations.

A useful contribution of this analysis is the identification of two distinct types of patterns that advance our understanding of how tensions in digital technology development manifest and how their resolutions can be conceptualised. The first, technology-driven resolutions, is about rendering trade-offs obsolete through innovative use of technology. The second, governance-driven, highlights the important role of transparent processes, user agency, and co-design in navigating socio-technical tensions that technology alone cannot solve.

For practitioners and policymakers, the message is: trade-off hacking is not about ignoring difficult choices but about confronting them with a mindset that places human values at the center of innovation. While the perfect win-win is seldom achievable, this approach is fundamental to the continuous process of building trust in digital technology, which is the essential ingredient for any digital transformation to succeed.

Ethics and consent

Ethical approval and consent were not required for this study.

Declarations

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author used ChatGPT (GPT-5) for minor improvements on language and readability. After using the tool, the author reviewed and edited the content and takes full responsibility for the content of the published article.

Data availability

No data associated are with the article. No datasets were generated or analyzed during the preparation of this case study.

References

Attoresi M: **Neuro-symbolic Artificial Intelligence.** *European data protection supervisor* 2025.
Reference Source .

Awad NF, Krishnan MS: **The personalisation privacy paradox: an empirical evaluation of information transparency and the willingness to be profiled online for personalisation.** *MIS Quarterly.* 2006; **30**(1): 13–28.
Publisher Full Text

Chellappa RK, Sin RG: **Personalization versus privacy: an empirical examination of the online consumer's dilemma.** *Inf Technol Managt.* 2005; **6**(2–3): 181–202.
Publisher Full Text

Cloarec J: **The personalisation-privacy paradox in the attention economy.** *Technol Forecast Soc Change.* 2020; **161**: 120299.
Publisher Full Text

Crook B, Schlüter M, Speith T: **Revisiting the performance-explainability trade-off in Explainable Artificial Intelligence (XAI).** *2023 IEEE 31st International Requirements Engineering Conference Workshops (REW).* 2023: 316–324.
Publisher Full Text

Detopoulou P, Voulgaridou G, Moschos P, *et al.*: **Artificial Intelligence, nutrition, and ethical issues: a mini-review.** *Clin Nutr Open Sci.* 2023; **50**: 46–56.
Publisher Full Text

Donovan SM, Abrahams M, Anthony JC: **Perspective: challenges for personalized nutrition in the current United States regulatory framework and future opportunities.** *Adv Nutr.* 2025; **16**(3): 100382.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

DRG4FOOD: *About. European Union Horizon Europe project.* n.d.
Reference Source .

Edelman: *Edelman trust barometer.* 2024, 2024.
Reference Source .

Ellahi RM, Wood LC, Bekhit AEA: **Blockchain-based frameworks for food traceability: a systematic review.** *Foods.* 2023; **12**(16): 3026.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Ferroli F, Hermsen K, Talmoudi K: **Governance concept for digital responsibility.** *Zenodo.* 2023.
Publisher Full Text

Ferroli F, Weich S: **DRG4FOOD digital responsibility report 2023/2024.** *DRG4FOOD Project* 2025.
Reference Source

Friedman B, Kahn PH, Borning A: **Value sensitive design and information systems.** *In: The Handbook of Information and Computer Ethics.* 2008: 69–101.
Publisher Full Text

Garimella KK, Conway D: **Zero-knowledge proofs and privacy: a technical look at privacy.** *In: Technology, work and globalization.* 2024: 157–179.
Publisher Full Text

Hassija V, Chamola V, Mahapatra A, *et al.*: **Interpreting black-box models: a review on explainable Artificial Intelligence.** *Cognit Comput.* 2023; **16**(1): 45–74.
Publisher Full Text

Identity Valley: *Digital responsibility goals.* 2025; Accessed September 10, 2025.
Reference Source

Jacobs D, McDaniel T: **A survey of user experience in usable security and privacy research.** *In: Lecture notes in computer science.* 2022: 154–172.
Publisher Full Text

Lennartsson M, Kävrestad J, Nohlberg M: **Exploring the meaning of usable security - a literature review.** *Inf Comput Secur.* 2021; **29**(4): 647–663.
Publisher Full Text

Loftus TJ, Ruppert MM, Shickel B, *et al.*: **Federated learning for preserving data privacy in collaborative healthcare research.** *Digit. Health.* 2022; **8**: 205520762211344.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Meier JJ, Hermsen K, Bauer J, *et al.*: **Digital responsibility goals – A framework for a human-centered sustainable digital economy with a focus on trusted digital solutions.** *Stud Health Technol Inform.* 2022; **293**: 250–259.
PubMed Abstract | **Publisher Full Text**

Mills AJ, Durepos G, Wiebe E: *Encyclopedia of Case Study Research.* 2010.
Publisher Full Text

Reis AA, Malpani R, Vayena E, *et al.*: **Ethics and governance of Artificial Intelligence for health: WHO guidance.** *World Health Organization* 2021;
Reference Source .

Sadek M, Mougenot C: **Challenges in value-sensitive AI design: insights from AI practitioner interviews.** *Int J Hum Comput Interact.* 2024; **41**(17): 10877–10894.
Publisher Full Text

Sanders EBN, Stappers PJ: **Co-creation and the new landscapes of design.** *CoDesign.* 2008; **4**(1): 5–18.
Publisher Full Text

Sedlmeir J, Smethurst R, Rieger A, *et al.*: **Digital identities and verifiable credentials.** *Business & Information Systems Engineering.* 2021; **63**(5): 603–613.
Publisher Full Text

Spiekermann S, Winkler T: **Value-based engineering with IEEE 7000TM.** *SSRN Electronic Journal.* 2022;
Publisher Full Text

Steinke J, Ortiz-Crespo B, Van Etten J, *et al.*: **Participatory design of digital innovation in agricultural research-for-development: insights from practice.** *Agric Syst.* 2021; **195**: 103313.
Publisher Full Text

Van Der Veer SN, Riste L, Cheraghi-Sohi S, *et al.*: **Trading off accuracy and explainability in AI decision-making: findings from 2 citizens' juries.** *J Am Med Inform Assoc.* 2021; **28**(10): 2128–2138.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Wu X, Oniani D, Shao Z, *et al.*: **A scoping review of Artificial Intelligence for precision nutrition.** *Adv Nutr.* 2025; **16**(4): 100398.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Xu H, Luo X, Carroll JM, *et al.*: **The personalization privacy paradox: an exploratory study of decision-making process for location-aware marketing.** *Decis Support Syst.* 2011; **51**(1): 42–52.
Publisher Full Text

Yu D, Yang B, Liu D, *et al.*: **A survey on neural-symbolic learning systems.** *Neural Netw.* 2023; **166**: 105–126.
PubMed Abstract | **Publisher Full Text**

Open Peer Review

Current Peer Review Status: ? ✓ ?

Version 2

Reviewer Report 13 March 2026

https://doi.org/10.21956/openreseurope.25084.r70356

© **2026 Kurdyś-Kujawska A.** This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

?

Agnieszka Kurdyś-Kujawska
Koszalin University of Technology, Kwiatkowskiego, Koszalin, Poland

General comments: The article is a coherent and valuable contribution to the discussion on digital ethics in the agri-food sector. It effectively combines a high level of theory with practical examples of implementation, making it useful for both researchers and software developers. A strong point of the study is its departure from theorizing about ethics in favor of showing how values can become the basis for technical creativity.
Abstract: The abstract is complete and contains all the key elements required of a scientific abstract.
Introduction: The introduction is well structured, with the author moving smoothly from the broad technological context to specific ethical dilemmas, and then to a concrete methodology and analytical framework.
Background: The background to the article is a well-constructed analytical tool that allows the reader to understand why and how the pilot projects studied attempted to overcome traditional technological limitations.
Hacking compromises: excerpts from the DRG4FOOD pilot: The section is structured around three key dilemmas, allowing for a clear comparison of different approaches to the same problem. The author does not limit himself to describing the projects, but uses them as “micro-case studies” to illustrate specific strategies. The selection of eight pilots ensures adequate technological diversity – from mobile applications to blockchain systems in the supply chain. The text accurately identifies the spectrum of interventions, which is one of the strongest points of this section. The author maintains the necessary objectivity, avoiding claims that all dilemmas have been completely eliminated. Despite its high analytical value, the text is based mainly on qualitative and narrative descriptions. It lacks hard data or harmonized indicators (metrics) that would allow for an objective comparison of the effectiveness of individual “hacks.” The analysis focuses on the declared strategies from the incubation period, which, without long-term research, does not provide complete certainty as to the ultimate effectiveness of these solutions when used on a large scale.
Synthesis: Two patterns of trade-off hacking: The section entitled “Synthesis: Two patterns of trade-off hacking” provides a theoretical and practical summary of the analysis, giving the work a systematic character. Here, the author moves away from describing individual projects in favor of

creating a universal framework for classifying strategies for resolving project conflicts. The analysis remains at the qualitative level. Although the author has created a solid theoretical framework, it lacks more precise criteria for choosing between a technological and a managerial pattern depending on, for example, the size of the organization or the degree of technological maturity of the project.
Discussion: This section constitutes the analytical layer of the work, going beyond the description of specific cases towards a synthesis of conclusions for the entire food technology sector. Despite the high quality of the text, I see areas that could be explored in greater depth to enhance the scientific value of this section. The discussion remains at a narrative and qualitative level. There is a lack of specific indicators that a developer could use to measure how much better a "hacked" compromise is than a standard solution. The author mentions the risk of "bottlenecks" (e.g., in human verification), but the discussion lacks a broader analysis of how solutions such as human-in-the-loop or co-design will fare in the transition from pilot phase to mass market scale. Although "low-margin markets" are mentioned, the discussion could have analyzed in more detail how small and medium-sized enterprises (SMEs) operating outside EU funds can realistically finance costly technological "hacks" (e.g., ZKP cryptography). The argument is very strongly rooted in the European reality. It would be an improvement to address whether these models are applicable in jurisdictions with different data protection standards or lower trust in institutions.
Conclusions: The conclusions section is a logical conclusion to the article. It effectively combines a high level of theory (digital ethics) with practical conclusions for developers and decision-makers, emphasizing that building trust in technology is a continuous and demanding process. Despite its high quality, there are some areas in the conclusions that could be explored in greater depth. The author mentions the costs in terms of complexity, but does not indicate how organizations should assess whether a given "hack" is economically viable in the long term. The conclusions focus on paths to success. There is a lack of emphasis on the boundary conditions that make trade-off hacking doomed to failure, which would be a valuable warning for practitioners.

Is the background of the case's history and progression described in sufficient detail?
Yes

Is the work clearly and accurately presented and does it cite the current literature?
Yes

If applicable, is the statistical analysis and its interpretation appropriate?
Not applicable

Are all the source data underlying the results available to ensure full reproducibility?
Partly

Are the conclusions drawn adequately supported by the results?
Yes

Is the case presented with sufficient detail to be useful for teaching or other practitioners?
Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: agricultural economics, risk management in agriculture, financing of the agri-food sector

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Reviewer Report 25 February 2026

https://doi.org/10.21956/openreseurope.25084.r70274

© 2026 Vargas-Canales J. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Juan Manuel Vargas-Canales
Universidad de Guanajuato, Salvatierra, Mexico

After reviewing the revised version of the manuscript, I can confirm that the previously raised observations have been satisfactorily addressed. The author implemented substantial and appropriate adjustments in each of the highlighted sections, strengthening the conceptual clarity, methodological transparency, and critical balance of the analysis.
In particular:

- The abstract now explicitly describes the analytical approach used.
- The introduction adopts a more neutral and analytical tone.
- The conceptual framework adequately expands on the notion of *trade-off hacking* and justifies the selection of the three trade-offs.
- The vignette section incorporates additional critical elements, including consideration of tensions that remain unresolved.
- The synthesis section deepens the discussion of limitations associated with the identified patterns.
- The discussion provides a more critical analysis regarding the project context, replicability, and external factors.
- The conclusions have been adjusted to reflect a more cautious and balanced tone.

With these modifications, the manuscript now presents a more robust, critical, and methodologically transparent version, satisfactorily addressing the recommendations previously provided.

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Economy, innovation, technological change, digital agriculture

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Version 1

Reviewer Report 03 February 2026

https://doi.org/10.21956/openreseurope.23550.r68250

© **2026 Vargas-Canales J.** This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

?

Juan Manuel Vargas-Canales
Universidad de Guanajuato, Salvatierra, Mexico

Trade-off hacking: Have the cake and eat it too? Turning competing interests into win-wins in data-driven food technology

General comments:

- The work addresses a real and highly topical problem. It is highly relevant to policy design. However, there are some areas that need improvement.

Abstract

- It is not specified how the pilot projects were explicitly analyzed.

Introduction

- This section contains statements heavily biased towards justifying the project. It would be preferable for it to be more neutral.

Background

- Trade-off hacking: this concept should be expanded upon as much as possible. It would be beneficial to explore the issues further.
- Regarding the trade-offs, it would be helpful to explain why there were only three.

Trade-off hacking: vignettes from the DRG4FOOD pilots

- This section would be strengthened if it were more critical. Furthermore, it would be advisable to include failures, unresolved tensions, and aspects not yet considered.

Synthesis: Two patterns of trade-off hacking

- In this section, it would be advisable to delve deeper into the limitations of each of the patterns. It should be more critical.

Discussion

- This section seems too closely aligned with the project objectives. A more critical analysis is needed; otherwise, there is a risk of analytical bias. What would happen without funding? What would happen if an attempt were made to replicate the study? What aspects are most important to consider? What happens outside the contexts of the analysis?

Conclusions

- This section is very encouraging. It shows a positive bias. It needs to be more critical.

Is the background of the case's history and progression described in sufficient detail?
Partly

Is the work clearly and accurately presented and does it cite the current literature?
Yes

If applicable, is the statistical analysis and its interpretation appropriate?
Not applicable

Are all the source data underlying the results available to ensure full reproducibility?
No source data required

Are the conclusions drawn adequately supported by the results?
Partly

Is the case presented with sufficient detail to be useful for teaching or other practitioners?
Partly

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Economy, innovation, technological change, digital agriculture

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Author Response 13 Feb 2026

Ferdinand Ferroli

Thank you for the constructive feedback. The manuscript has been revised to improve clarity, balance, and critical reflection. **Abstract**
The description of how pilot material was analysed has been clarified by stating the overall approach, data basis, and main analytical process more directly. **Introduction**
The tone has been revised to be more neutral and analytical, with less project-affirming language. **Background (trade-off hacking and trade-off selection)**
The conceptual discussion of trade-off hacking has been expanded, and clearer rationale has been added for focusing on the three main trade-offs in scope. **Pilot vignette section**
A stronger critical dimension has been added through discussion of unresolved tensions and aspects that may remain challenging in practice. **Synthesis section**
More explicit reflection has been included on the limitations and conditions of the identified patterns. **Discussion**
The critical analysis has been deepened, including reflections on likely constraints in other

implementation settings. **Conclusions**
The conclusion has been revised to a more cautious framing, with clearer acknowledgement of limits and open questions.

Competing Interests: No competing interests were disclosed.

Reviewer Report 03 February 2026

https://doi.org/10.21956/openreseurope.23550.r65045

© 2026 Galanakis C. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

?

Charis M Galanakis
Galanakis Laboratories, Chania, Greece

This study aims to reconceptualise ethical and functional trade-offs in data-driven food technologies through the lens of “trade-off hacking,” shifting from zero-sum compromises to innovation-oriented, user-centric solutions. By analysing eight Horizon Europe DRG4FOOD pilot projects, it identifies recurring tensions (privacy versus personalisation, performance versus explainability, and security versus user experience). It systematically examines how they are addressed in practice. The main outcome is a typology of reproducible design patterns that distinguishes technology-driven from governance-driven resolutions and demonstrates how trust, responsibility, and performance can be mutually reinforcing. The study is conceptually original, empirically grounded, and offers a compelling, well-argued contribution to responsible digital innovation and is highly commendable.

1. Clarify the methodological approach more explicitly, including case selection criteria, data sources, and analytical steps, to strengthen transparency and replicability and to help readers better assess the robustness and transferability of the findings across contexts.
2. Provide a more apparent justification for focusing on only three trade-offs, discussing whether other relevant tensions were considered or excluded, and how this choice may shape interpretations and limit the generalisability of conclusions in broader food systems.
3. Expand the discussion of limitations by more critically reflecting on potential biases introduced by the responsible incubator setting, funding incentives, and coaching support, and how these factors might affect applicability outside Horizon Europe programmes and markets.
4. Include more concrete comparative analysis across the pilot cases, possibly through a synthesis table or metrics, to better demonstrate the relative effectiveness of different strategies and avoid an overly descriptive, narrative-driven assessment of outcomes and impacts observed.
5. Add clearer implications for practitioners by translating identified patterns into actionable design guidance or decision criteria, helping developers understand when to prioritise technological solutions versus governance-driven approaches under real-world constraints, including scale, cost, and maturity levels.

Is the background of the case's history and progression described in sufficient detail?
Yes

Is the work clearly and accurately presented and does it cite the current literature?
Yes

If applicable, is the statistical analysis and its interpretation appropriate?
Not applicable

Are all the source data underlying the results available to ensure full reproducibility?
Partly

Are the conclusions drawn adequately supported by the results?
Partly

Is the case presented with sufficient detail to be useful for teaching or other practitioners?
Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: food security; food waste

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Author Response 13 Feb 2026

Ferdinand Ferroli

Thank you for the positive and helpful review. The manuscript has been revised in line with the main suggestions.

1. **Methodological transparency**
   The methodological approach is now presented more explicitly, including case selection, evidence sources, and analytical steps.
2. **Why three trade-offs**
   Clearer justification has been added for focusing on the three selected trade-offs, and the scope implications of this choice are now acknowledged more directly.
3. **Limitations and transferability**
   The limitations discussion has been expanded, with more explicit reflection on contextual influences and on transferability beyond similar programme environments.
4. **Cross-case comparison**
   Comparison remains qualitative in this version; no harmonised effectiveness metrics were introduced due to the exploratory design and heterogeneous evidence base. This is now stated as a limitation and identified as a priority for future research.

5. **Implications for practice**
   Practitioner-oriented guidance has been refined to make the identified patterns more actionable under real-world constraints.

Competing Interests: No competing interests were disclosed.