From principle to practice: a strategic framework for responsible AI in the FMCG sector

TABLE OF CONTENTS

A practical guide to building a responsible AI framework

This document provides a structured framework for developing a responsible AI program, with a specific focus on offering a scalable model that can be adapted by small and mid-sized enterprises (SMEs) and larger corporations. The goal is to move beyond theory and provide a practical toolkit for building effective, value-driven governance.

1. Foundational pillars: core principles for companies working within the FMCG ecosystem

Based on a review of common practices, the following points are a solid start for any FMCG company looking to create its own policy. They show a shared understanding of key risks and ethical issues in the sector, which can help improve efficiency and compatibility between systems. It is important not only to list these principles but to articulate their specific relevance to the FMCG context.

- Accountability: Clear ownership must be assigned for every AI system throughout its lifecycle. In FMCG, this means the brand manager for a marketing AI, the supply chain director for a forecasting model, and the R&D lead for a product formulation algorithm are all explicitly responsible for the system's ethical performance and outcomes. Accountability ensures that if an issue arises, there is a clear line of sight to the individual empowered to address it.

- Transparency & explainability: Stakeholders (consumers, employees, regulators) should be informed when they are interacting with an AI system, and the company must be able to explain, in understandable terms, how an AI system arrived at a particular decision. For an FMCG company, this could mean disclosing the use of an AI-powered chatbot for customer service or being able to explain to a customer why a pricing algorithm sets a specific price in a particular market. This builds trust and is essential for debugging and auditing.

- Fairness: AI systems must be designed and evaluated to prevent unfair bias and discriminatory outcomes. In FMCG, this is not an abstract concept. An AI model for promotion targeting that unfairly excludes certain demographic groups is not only unethical but represents a lost market opportunity. A biased AI for supply chain forecasting can create stockouts in specific neighbourhoods, directly impacting sales and brand reputation. Proactive bias mitigation is both an ethical duty and a

commercial necessity.

- Security & reliability: AI systems, and the data they rely on, must be protected from theft, corruption, and unauthorised access. They must also be robust, performing consistently and accurately under a variety of conditions. For an FMCG company, a compromised pricing algorithm could lead to massive financial loss, while an unreliable demand forecasting model could cripple the supply chain. Technical integrity is the foundation upon which all other principles rest.

- Privacy & data governance: The collection, use, and storage of data, particularly personal data from consumers, must adhere to the highest standards of privacy and ethical governance. FMCG companies are custodians of vast amounts of consumer data from loyalty programs, e-commerce sites, and marketing campaigns. Misusing this data or failing to protect it can cause irreparable damage to consumer trust and brand loyalty.

- Human-in-the-loop: Critical decisions should always have a meaningful level of human oversight. While AI can recommend, analyse, and predict, the final accountability for high-stakes decisions, such as a major product recall identified by an AI quality control system or significant changes to workforce allocation based on an efficiency model, must rest with a human.

- Requirement for agentic systems: As systems become more agentic (autonomous), organisations may consider implementing specific review checkpoints. This includes evaluating the scope of autonomy and ensuring clear "Human-in-the-Loop" intervention capabilities.

2. The governance engine: structuring for accountability and oversight

Principles are meaningless without a structure to enforce them. The *"governance engine"* is the set of roles, responsibilities, and processes that brings a policy to life. The complexity of this engine should scale with the size and AI maturity of the organisation. Key roles include:

- AI project owner: The business leader sponsoring the AI initiative. This individual is responsible for defining the business case, ensuring the project aligns with company principles, and is accountable for the system's performance and impact.
- Technical reviewer(s): Data scientists, engineers, or IT security specialists responsible for assessing the technical aspects of an AI system. They evaluate its robustness, security, data integrity, and the technical methods used for bias detection.
- Ethics/compliance reviewer(s): Representatives from legal, compliance, or a dedicated ethics function. They are responsible for assessing the AI project's alignment with the company's principles, relevant laws and regulations, and potential societal impact.
- Executive sponsor / governance body: For low-risk projects, this may be a department head. For high-risk or enterprise-wide projects, this should be a formal, cross-functional committee or board (like the central governance body model) responsible for providing final approval and executive oversight.

3. From policy to practice: an implementation roadmap for small and mid-sized enterprises

For SMEs, the prospect of building a comprehensive AI governance program can be intimidating. The key is to start with a lightweight, practical framework that can mature over time. The following five-step roadmap provides a clear path forward.

- Step 1: adopt and communicate principles: Begin by formally adopting the foundational pillars from section 3.1. Write them in a simple one-page document. Communicate this document to all employees to establish a shared understanding of the company's commitment to responsible AI.

- Step 2: establish "lightweight" governance: Implement a simple but formal approval process. This can be modelled on the decentralised approach: all new uses of AI tools or the initiation of any AI project require a documented proposal and written approval from a designated manager or a small, standing committee of two or three leaders (e.g., heads of IT, Marketing, and Operations). This creates a crucial checkpoint without requiring a large bureaucracy.

- Step 3: implement a "prime directive" for data security: The single greatest immediate risk for many SMEs is the leakage of confidential data into public AI tools. Adopt a clear, simple, and non-negotiable rule inspired by the most risk-averse policies: No confidential company information is to be entered into any public, non-approved AI tool. This includes customer lists, financial data, strategic plans, proprietary formulas, and internal employee information.

- Step 4: create simple usage guidelines: Provide employees with practical, easy-to-understand "do's and don'ts." This should include advice modelled on best practices, such as using AI as a "guide" for inspiration and first drafts but not as a replacement for human work; always verifying the accuracy of AI-generated facts; and being transparent with colleagues when AI has been used to assist in creating work.

- Step 5: scrutinise your vendors: SMEs often rely heavily on third-party software for functions like marketing automation, CRM, or logistics. When procuring any new software that is advertised as "AI-powered," make vendor scrutiny a standard part of the due diligence process. Ask contracts to include requirements for ongoing monitoring of live systems to detect model drift or emergent issues post-deployment. Ask vendors to provide their own responsible AI policy and inquire about their data handling practices and model transparency, echoing the

requirements of more mature corporate policies. This is a matter of individual company due diligence, not a collective assessment of suppliers.

This roadmap is designed to be scalable. As an SME grows and its use of AI becomes more sophisticated, it can progressively add more layers of rigor. The manager-approval process can evolve into a formal committee, simple usage guidelines can become a comprehensive training program, and vendor scrutiny can become a full third-party risk management framework.

The following table illustrates this scalable model.

Governance domain	Basic (getting started)	Intermediate (building capability)	Advanced (strategic advantage)
Principles	Adopt and communicate a core set of 5-6 principles.	Integrate principles into project management and employee onboarding.	Publish a public-facing Responsible AI Pledge or Report.
Governance body	Designated Manager or small committee provides approval.	Formal, cross-functional review committee with defined charter.	Dedicated AI Centre of Excellence or executive-level central governance body.
Risk assessment	Informal checklist reviewed by manager/committee	Standardised Risk Assessment template required for all new projects.	Tiered risk assessment process: high-risk projects require in-depth ethical review.
Employee guidelines	Simple "Do's and Don'ts" document.	Mandatory annual training on responsible AI policy and procedures.	Role-specific training; cultivation of "AI Ethics Champions" within business units.
Vendor Management	Ask vendors for their AI policy during procurement.	Require AI-specific clauses in vendor contracts; conduct third-party risk assessments.	Proactive auditing of key AI vendors; sharing best practices on ethical AI and data compliance.

4. Advanced strategies: futureproofing your AI governance

For companies aiming for the "Advanced" tier of maturity, governance must evolve from a static control function to a dynamic, strategic capability. This involves embracing the forward-looking practices that currently represent gaps in most corporate policies.

- Implementing adaptive governance: The AI landscape is not static, and neither should be your governance. Establish a formal process for reviewing and updating the AI policy on at least an annual basis, or more frequently in response to major technological or regulatory shifts. Create an internal "red team" or engage third-party experts to proactively test models for bias, security flaws, and other potential harms before they are deployed. Implement post-deployment monitoring systems to track model performance and detect "concept drift" or emergent biases over time, ensuring the system remains fair and reliable long after its initial launch.

- Cultivating a responsible AI culture: Move beyond rules-based training to encourage a genuine culture of critical thinking and ethical inquiry. This involves creating psychologically safe channels such as a confidential reporting hotline model for employees to raise concerns about AI systems without fear of retaliation. It also means actively celebrating and rewarding examples of responsible innovation. The goal is to shift the mindset from "am I allowed to do this?" to "is this the right thing to do?" This cultural shift, which addresses employee concerns and focuses on enablement, is a powerful competitive differentiator.

- Building trust through radical transparency: The most advanced companies understand that trust is a tangible asset. Go beyond internal compliance by creating external-facing transparency reports, like sustainability reports, which detail the company's approach to responsible AI, the governance structures in place, and the progress being made. Consider developing a formal, public "Responsible AI Pledge" to clearly articulate your commitments to consumers, partners, and the public. These actions transform governance from an internal process into a powerful brand-building tool.

By adopting these advanced strategies, FMCG companies can move beyond mere risk mitigation. They can position themselves as leaders in the ethical application of technology, building a foundation of trust and resilience that will be essential for navigating the complexities of the AI-driven future.

Appendix: A modular policy toolkit

This appendix provides a set of practical, "plug-and-play" resources to help companies, particularly SMEs, get started on the development of their responsible AI framework.

These templates can be adapted to fit the specific needs and context of your organisation:

- Template clauses for a responsible AI policy
- AI project intake list
- Vendor due diligence questionnaire for AI systems

Template clauses for a responsible AI policy

Use these pre-written text blocks as a starting point for drafting your own policy document.

1. Scope & applicability:

> "This Policy applies to all employees, contractors, and other personnel (collectively, 'Associates') of [Company Name] and its consolidated subsidiaries. It governs the design, development, procurement, deployment, and use of all Artificial Intelligence (AI) Systems within the company. Third parties developing or providing AI Systems to or on behalf of [Company Name] must comply with this Policy or one that is substantially similar as a condition of their engagement."

2. Core principles:

> "[Company Name] is committed to the responsible and ethical use of AI.
>
> All AI activities must adhere to our core principles:
>
> - Accountability: We assign clear ownership for our AI systems and their outcomes.
> - Transparency: We are open and clear about using AI with our stakeholders.
> - Fairness: We work diligently to identify and mitigate unfair bias in our AI Systems.
> - Security & reliability: We ensure our AI Systems are secure, robust,

and perform as intended.

- Privacy: We respect individuals' privacy and protect their data.
- Human oversight: We ensure meaningful human oversight for critical decisions."

3. Third-party requirements:

"When procuring AI Systems or engaging vendors who use AI in the delivery of their services, [Company Name] will conduct due diligence to assess their commitment to responsible AI.

All relevant contracts must include provisions requiring the vendor to adhere to principles consistent with this Policy, particularly concerning data privacy, security, and compliance with applicable laws."

AI project intake checklist

Before any new AI project is approved, the project owner should be required to complete the following checklist. This ensures that key ethical and risk considerations are addressed from the outset.

Question	Response (Yes/No/Details)
1. Purpose & value	
Is the business objective of this AI system clearly defined?	
Does the project align with our company's values and Code of Conduct?	
2. Data & privacy	
Will this system use Personal Data (as defined in our Privacy Policy)?	
If yes, has a Privacy Impact Assessment been	

completed with the Privacy Team?	
Is the data to be used representative of the population it will affect?	
3. Fairness & bias	
Is there a risk that this system could produce biased or discriminatory outcomes for any group of individuals?	
What specific steps will be taken to test for and mitigate bias in the data, model, and outcomes?	
4. Transparency & oversight	
How will users or affected individuals be made aware they are interacting with an AI system?	
How will the system's decisions be validated for accuracy?	
What is the process for human intervention or appeal if the system makes an error?	
5. Security & reliability	
Has the system been assessed for potential security vulnerabilities?	
What is the plan for monitoring the system's performance after deployment?	

Vendor due diligence questionnaire for AI systems

When evaluating a third-party AI solution or vendor, use these questions to assess their ethical posture and risk profile.

1. Can you provide a copy of your organisation's 'Responsible or Ethical AI' policy?
2. How do you test and mitigate bias in your AI models? Can you provide documentation on your methodology?
3. What level of transparency do you provide into your model's decision-making process? Is it possible to get an explanation for a specific outcome?
4. Where will our company's data be stored, and what security measures are in place to protect it?
5. Will our company's data be used to train your models for other customers? If so, is there an option to opt-out?
6. How does your system help us ensure compliance with relevant data protection regulations?
7. What processes do you have in place for data retention and deletion upon termination of our contract?
8. Can you describe the human oversight mechanisms built into your system or recommended for its use?
9. Does your organisation provide documented protocols for the ongoing monitoring of live systems to ensure performance remains aligned with ethical standards?