Foresight Canvas / Futures Wheel Guide

A speculative design process to anticipate the long-term, unintended consequences of our agent. This audit focuses on identifying second-order effects, potential for misuse, and systemic risks.

Applying the Foresight & Feedback Loop to your agent helps to stress-test against future scenarios, surface unintended consequences, and ensure long-term viability and trust in an evolving world.

Instructions

01 Define the core innovation or capability your agent introduces at the center of the wheel.

02 Identify first-order consequences: direct, immediate effects on users, workflows, and stakeholders.

03 Map second-order consequences: ripple effects that emerge from those direct impacts, especially on organizations, markets, and society.

04 Continue mapping higher-order effects as they branch outward, paying attention to unintended or non-obvious outcomes.

05 For each consequence, assess whether it is positive, neutral, or negative and consider who is affected.

06 Highlight areas of uncertainty or risk that require deeper investigation, monitoring, or design intervention.

Methodology

The Futures Wheel is a structured brainstorming tool originally developed by Jerome Glenn. It starts with a central event or change and maps outward in concentric rings of consequences. In the context of agent design, it helps teams move beyond immediate functionality to explore the broader systemic effects of deploying autonomous agents. By systematically tracing cause-and-effect chains, designers can surface hidden risks, anticipate unintended behaviors, and identify opportunities for proactive safeguards. This method is especially valuable in speculative design settings where the goal is to stress-test agent capabilities against plausible future scenarios before they reach production.

Result: A Futures Wheel artifact documenting all orders of consequences related to the agent's core innovation, usable as a living reference for design decisions, risk assessments, and stakeholder communication.

The Futures Wheel



Reading the wheel. Start with the problem and the tech solution at the center. Work outward to first- and second-order consequences, considering users, security, and your company. Add further stakeholders and higher-order effects as needed.

Worksheet adaptation: the blank wheel retains the source's concentric structure, stakeholder prompts, and consequence levels. The register adds room for the assessment requested in the instructions.

Adapted from Outshift Design. Source content and diagrams © 2026 Outshift.

Foresight Canvas / Futures Types Guide

Classify scenarios from the wheel and translate their implications into design priorities.

Instructions

01 Start from the consequences identified in the Futures Wheel.

02 Classify each scenario into one of four futures types: Positive, Negative, Plausible, or Dystopian.

03 For each classification, describe the scenario and its implications for users, systems, and society.

04 Identify design responses: what guardrails, fallback mechanisms, or ethical boundaries can be embedded to steer toward positive futures and mitigate negative ones?

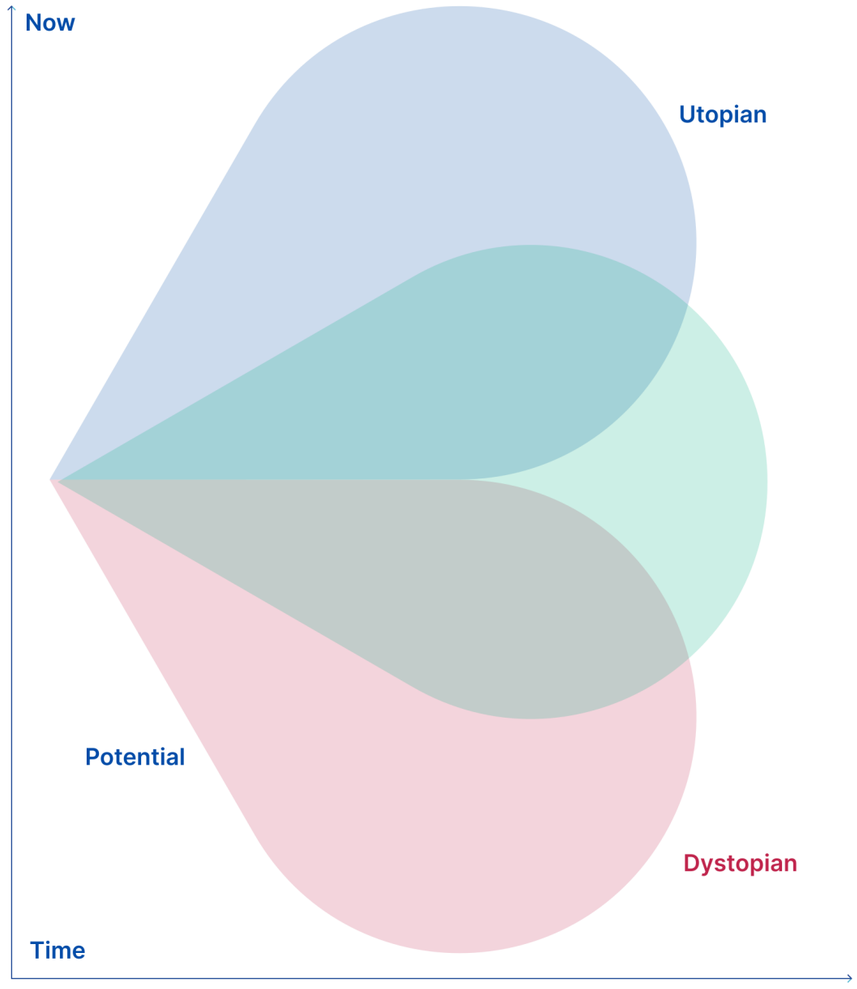
05 Use the resulting map to prioritize design interventions based on likelihood and severity of impact.

Methodology

The Futures Types framework draws on established futures studies methodology, particularly the work of Jim Dator and Sohail Inayatullah. It categorizes possible futures into archetypes, helping designers move from abstract speculation to actionable design priorities. By placing agent-related scenarios into Positive, Negative, Plausible, and Dystopian quadrants, teams gain a structured way to evaluate not just what could happen, but how desirable or dangerous each outcome might be. This approach bridges the gap between speculative foresight and practical design decisions, ensuring that autonomous systems are built with resilience, accountability, and ethical awareness from the start.

Result: A Futures Types map that classifies agent-related scenarios into four quadrants, providing a clear framework for prioritizing design interventions and embedding safeguards against the most impactful negative outcomes.

The futures illustration



From illustration to worksheet. The source image shows a cone of futures labeled Now, Utopian, Potential, Dystopian, and Time. The source instructions instead ask for four categories: Positive, Negative, Plausible, and Dystopian.

The worksheet uses those four requested categories as writable quadrants. These are discussion lenses: a plausible scenario can also be positive or negative. Use the priority row to compare likelihood and severity before choosing an intervention.

Adapted from Outshift Design. Source content and diagrams © 2026 Outshift.

Foresight Canvas / 01 Futures Wheel

Trace the consequences of the agent's core innovation: problem → tech solution → first-order effects → second-order effects and beyond.

Designed for / agent	*Designed by / team*	*Date*	*Version*

Map outward

Begin with the problem and core capability. Draw arrows from each cause to its consequences.

First order

What will likely happen when the tech solution is implemented?

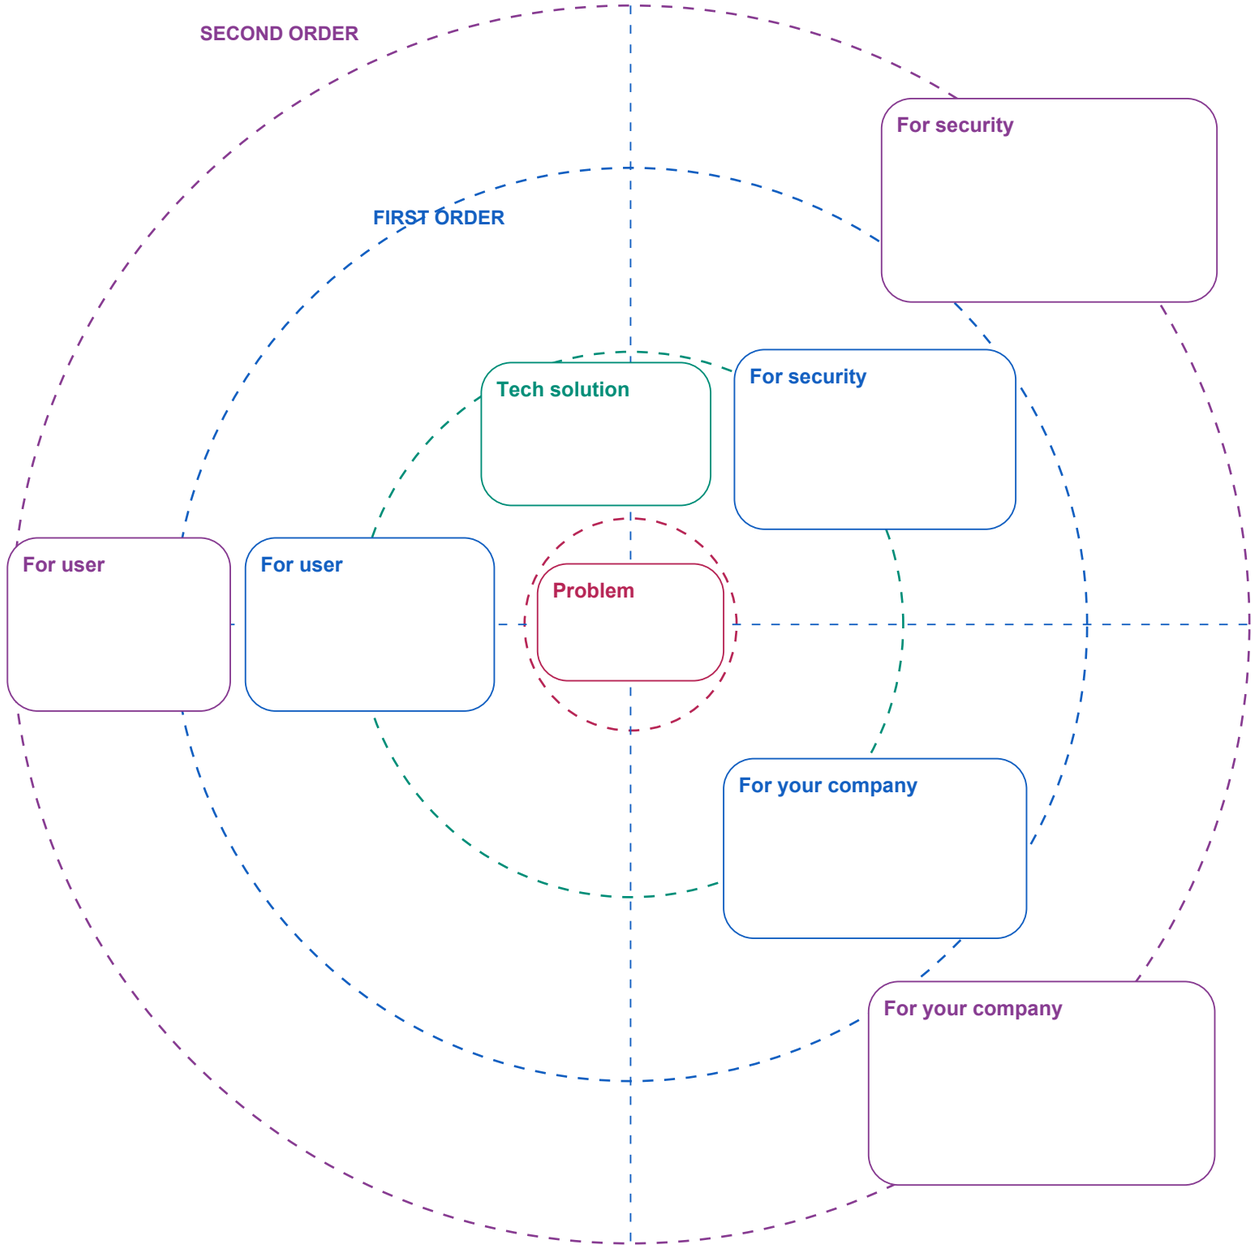
Second order

What will happen after the first-order consequence happens?

Assess every effect

Mark + positive, 0 neutral, or - negative. Note who is affected and flag uncertainty.

Extend the wheel with higher-order effects. Use IDs and arrows to preserve each cause-and-effect chain.



Consequence register

ID / who is affected / +, 0, -

Uncertainties & risks

What requires investigation, monitoring, or design intervention?

Higher-order effects

What happens next? Add linked IDs.

Adapted from Outshift Design. Source content and diagrams © 2026 Outshift.

Foresight Canvas / 02 Futures Types

Bring forward scenarios from the wheel. Describe their implications and choose guardrails, fallback mechanisms, and ethical boundaries.

Designed for / agent	*Designed by / team*	*Date*	*Version*

Positive

Scenario / source consequence IDs
Implications for users, systems, and society
Design response: guardrails, fallbacks, or ethical boundaries

Negative

Scenario / source consequence IDs
Implications for users, systems, and society
Design response: guardrails, fallbacks, or ethical boundaries

Plausible

Scenario / source consequence IDs
Implications for users, systems, and society
Design response: guardrails, fallbacks, or ethical boundaries

Dystopian

Scenario / source consequence IDs
Implications for users, systems, and society
Design response: guardrails, fallbacks, or ethical boundaries

Prioritize design interventions

Scenario / consequence IDs	Likelihood	Severity	Design intervention	Owner	Monitor / revisit

Adapted from Outshift Design. Source content and diagrams © 2026 Outshift.