

The Innovation Codex

Project Swift Worked Examples: Online Set

Appendix A tools A4, A6, A8, A10 to A13, and A15 to A17

Project Swift is a composite teaching case. Every Swift figure in this file is an illustrative composite, not a measurement.

This file holds the worked Project Swift examples for the ten tools whose examples appear online rather than in the book. The worked examples for A1 to A3, A5, A7, A9, and A14 appear in appendix A of the book. Each example shows the expected level of detail; read it before completing the blank form.

Contents

- A4. Team Charter Canvas
- A6. Portfolio View
- A8. Adoption Checklist
- A10. Risk Sandbox Charter
- A11. One-Page Story
- A12. Minimum Viable Adoption (MVA) Plan
- A13. Weekly Adoption Summary
- A15. Model Card
- A16. Innovation Performance Review
- A17. Future-Back Plan

Readers may reproduce the toolkit templates in appendix A, and those available at innovationcodex.com, for internal, noncommercial use, with attribution.

Version 1.1, October 2026.

A4. Team Charter Canvas

Book reference: appendix A4. Fillable form: innovationcodex.com/team-charter

Tool purpose. Establish or reset a team's working norms, roles, decision rights, and speaking-up practices before the work accelerates.

Worked example: Project Swift

Purpose. This canvas defines the working norms, roles, and communication plan for the Project Swift team so the pilot has clear decision rights, reliable communication, and explicit norms for surfacing bad news.

Team Mission. Reduce average counter time for returns to $\leq 3:30$ and improve the associate experience.

Scope. Validate the returns-counter problem with 20 target users, build the guided returns path, and run the 10-store pilot. Out of scope: automation and any rollout beyond the pilot.

Roles and Decision Rights

- **Operations lead (owner):** Swift rollout, store champions, performance monitoring. Decides: the scale, hold, or stop call, with Product, Support, and Data sign-off, and whether to renew the sandbox.
- **Product lead:** Guided path design, exception rules, decision cadence. Decides: changes to the guided path and the exception rules.
- **Data analyst:** Telemetry integrity, daily exports, dashboard health. Decides: whether the telemetry is reliable enough to support a decision.
- **Support lead:** Ticket taxonomy, support thresholds, and Swift support playbooks. Decides: changes to the ticket taxonomy and support playbooks.
- **Site leads:** On-floor coaching for Maria and the other returns-counter staff; collect friction notes. Decides: peak-shift coverage and coaching in their store.

Communication and Meeting Plan. Daily 10-minute stand-up and twice-weekly evidence review; dedicated Swift channel for pilot stores, weekly digest to stakeholders, Decision Briefs archived in the team drive.

Working Agreements. One decision owner per call, decisions recorded in a one-page brief.

Values in Action. Evidence over seniority: settle disagreements with data or the smallest test, not by rank. Learning over blame: hold a 30-minute retrospective after every failed experiment.

Norms That Protect Candor. "Our job is to learn, not to be right. Surfacing risk early is a valued contribution." Any team member may halt a review to surface a data concern.

Escalation Signals. Red if average counter time exceeds 4:00 for two days, or tickets exceed 8 per 1,000 uses for two days, or associate completion of the Swift primary path falls below 80 percent.

Definition of Done (Pilot). Activation ≥ 80 percent, average counter time $\leq 3:30$, tickets ≤ 8 per 1,000 uses, site leads confirmed and scheduled for peak shifts.

A6. Portfolio View

Book reference: appendix A6. Fillable form: innovationcodex.com/portfolio-view

Tool purpose. Provide the investment board a single, legible summary of every initiative and its stage.

Worked example: Project Swift

Purpose. This template provides the investment board with a clear portfolio-row summary of Project Swift's status, evidence, and next funding trigger.

Initiative. Project Swift (Retail Returns Flow).

Stage. Pilot to rollout.

Last Decision Date. Month 10 (Decision Brief: scale).

Next Trigger. Activation at or above 80 percent in live stores by Month 13, with average counter time $\leq 3:30$.

Funding Step Requested. Tranche 3 (staged regional rollout).

Evidence Snapshot. Activation 81 percent (returns completed without a manager), average counter time 3:18, customer completion at or above 80 percent, NPS plus 27, tickets 6 per 1,000 uses (pilot evidence dated Month 10).

Example portfolio row: Initiative: Project Swift, Retail Returns Flow. Stage: pilot to rollout. Last decision date: Month 10. Next trigger: activation at or above 80 percent in live stores by Month 13, with average counter time $\leq 3:30$. Funding step requested: Tranche 3. Evidence snapshot: activation 81 percent (returns completed without a manager), average counter time 3:18, NPS plus 27 (pilot sources dated Month 10).

- Primary owner: operations lead.
- Top risk: reversion to old workflow under peak pressure.
- Status: In progress.

A8. Adoption Checklist

Book reference: appendix A8. Fillable form: innovationcodex.com/adoption-checklist

Tool purpose. Make cross-functional launch readiness visible, trackable, and accountable.

Worked example: Project Swift

This checklist assessed readiness for the regional expansion from 50 to 150 stores. Each line carried an owner, a status, and the evidence that closed it.

Interface and workflow tested. Owner: product lead. Status: Done. Evidence: the guided path passes on the oldest terminal in the fleet.

Support trained and ticketing updated. Owner: support lead. Status: Done. Evidence: Swift ticket tags are live, and playbooks are published.

Communications and training ready. Owner: operations lead. Status: Done. Evidence: the five-minute walkthrough and one-page quick steps are distributed, and store champions are named.

Rollback plan defined and tested. Owner: product lead. Status: Done. Evidence: the fallback to the control flow is verified per store, with the data analyst confirming telemetry continuity.

Telemetry and dashboard live. Owner: data analyst. Status: Done. Evidence: daily exports and the adoption dashboard are refreshing.

Compliance sign-off. Owner: Legal. Status: Done. Evidence: exception rules are reviewed and codesigned with compliance.

A10. Risk Sandbox Charter

Book reference: appendix A10. Fillable form: innovationcodex.com/risk-sandbox

Tool purpose. Create a preapproved, time-bound space to run high-speed experiments without ad hoc sign-off.

Worked example: Project Swift

To test the guided returns path without per-store executive sign-off, the team worked inside a chartered sandbox.

Boundaries. 10 pilot stores, 14 days, the returns flow only. The sandbox may apply preapproved rules for the top 10 exception codes; all other cases fall back to the standard flow.

Preapprovals. Standing approval to run the guided path and export daily telemetry within the 10 stores, so no case-by-case sign-off is needed for each test day.

Automatic stop rules. Halt the guided path in any store if the defect rate rises 0.5 percentage points above baseline for two consecutive days, if customer completion drops below 80 percent, or if manager-override usage falls to near zero (1 percent or less of guided-path returns) for two consecutive days.

Data and privacy guardrails. Point-of-sale telemetry only; no personally identifiable customer data leaves the store systems.

Owner and expiry. The operations lead owns the sandbox; the charter expires at the end of the 14-day window and must be deliberately renewed to continue.

A11. One-Page Story

Book reference: appendix A11. Fillable form: innovationcodex.com/one-page-story

Tool purpose. Provide the team with one repeatable narrative that moves a specific audience to a decision.

Worked example: Project Swift

This One-Page Story was written for the board.

Audience. The board, deciding whether to fund Tranche 2.

Headline. A guided returns path gives associates back time and shortens the line.

Situation. At the returns counter, associates toggle between screens and a paper policy binder, often waiting for manager overrides, adding 60 to 120 seconds per return while customers wait.

What changed. A guided path with preapproved rules for common exceptions eliminated manual policy lookups in routine cases.

Evidence. In a 10-store pilot, average counter time fell to 3:18, 81 percent of returns were completed without a manager, NPS rose 27 points, and returns-related tickets fell from 18 to 6 per 1,000 uses.

The ask. Fund Tranche 2 (rollout readiness) for the 50-store launch.

Why now. Each week of delay leaves the queue friction and labor cost in place across every store.

A12. Minimum Viable Adoption (MVA) Plan

Book reference: appendix A12. Fillable form: innovationcodex.com/mvap

Tool purpose. Move an innovation from shipped to used by deliberately planning its adoption on one page.

Worked example: Project Swift

This MVA plan was designed to move the guided returns path from shipped to used.

Target audience. Returns-counter staff, starting with peak-shift staff such as Maria.

Adoption goal. At least 80 percent of eligible returns processed through the guided path within four weeks of store activation, not merely available on the screen.

Triggers. The guided path is the default on the returns screen, with a one-tap prompt when an eligible return begins.

Success metrics. Active usage rate (guided-path returns as a share of eligible returns), manager-override rate, and counter time, measured as use.

Habit-formation tactics. A named site lead coaches during peak shifts in week one; a five-minute walkthrough and one-page quick steps are posted at the counter; the Weekly Adoption Huddle reviews usage and removes friction.

Owner and cadence. The operations lead owns adoption; the Weekly Adoption Summary reports usage and blockers each week.

A13. Weekly Adoption Summary

Book reference: appendix A13. Fillable form: innovationcodex.com/weekly-adoption

Tool purpose. Provide the weekly huddle with a standard report that turns data into one committed change.

Worked example: Project Swift

This summary covered week one of the regional expansion from 50 to 150 stores.

Reporting period. Month 13, week 1.

Primary metric. Activation rate in the new regional stores 72 percent against an 80 percent goal. Status: Red.

Supporting metrics. Support tickets increased 10 percent in the new regions; counter time held steady; manager-override rate remained stable.

Top blocker. The scanners in the new stores fail on worn or faded receipts, pushing associates back to the manual path.

Change to test this cycle. Add a recovery path for unreadable receipts: an inline barcode scan hint and a manual-entry button (chapter 13).

Owner and next review. The product lead owns the fix; the change will be reviewed at next week's huddle.

A15. Model Card

Book reference: appendix A15. Fillable form: innovationcodex.com/model-card

Tool purpose. Document an AI model or automated decision system so its behavior and decision logic are transparent, auditable, and understandable to nonexperts.

Worked example: Project Swift

This Model Card documents Project Swift's returns preapproval system.

Intended use. Approve standard returns that fall within documented policy and route exceptions to a human. It assists the associate; it does not replace the associate's judgment.

Out-of-scope uses. Fraud adjudication, employee performance scoring, and any decision outside the documented return policy.

Performance. On the pilot set, the system handled the 10 highest-volume exception codes; approvals were audited against policy, and any case outside the rules was routed to a human.

Go-live criteria. Recommendations matched human decisions at 99 percent or higher on the validation set, with predefined thresholds for error cost, edge-case performance, and harm and fairness; low-confidence cases escalate to a person; drift is monitored against the validation baseline, with a tested rollback path.

Known limitations. The system depends on a readable receipt or label, does not handle edge cases not covered by the rules, and routes worn or faded receipts to the manual path.

Ethical considerations and safeguards. Associate performance data is used for system improvement and coaching, not punishment; rules are reviewed for bias using historical data; a manager override is always available; and near-zero override usage is treated as a user-agency warning signal.

Human-oversight path. Low-confidence cases and any case outside the rules escalate to a person; the visible Manager Override stays available on every transaction, and overrides are logged to find gaps in the rules.

Owner and review. The product lead owns the system; the card is reviewed whenever the rules, data, or intended use change.

A16. Innovation Performance Review

Book reference: appendix A16. Fillable form: innovationcodex.com/innovation-performance-review

Tool purpose. Assess contribution across four dimensions that feature delivery alone does not capture.

Worked example: Project Swift

Maria's review across the four dimensions.

Learning and de-risking. Proposed a two-day concierge test that manually delivered three metrics to the store manager, showing that he got the value he needed in 30 seconds and avoiding a multisprint dashboard build. This Wise Kill counts as a contribution, not a nondelivery.

Collaboration and influence. Moved from the subject of observation to an active partner, challenging proposed features in the team's own language: "What problem does this solve?" She reshaped the roadmap by making the user's work visible.

Ethical responsibility. Advocated for designs that preserve associate agency and flagged when a guided path risked removing judgment on legitimate edge cases.

Stewardship. Protected scarce engineering capacity by replacing an assumed build with a simple daily export of three numbers.

A17. Future-Back Plan

Book reference: appendix A17. Fillable form: innovationcodex.com/future-back-plan

Tool purpose. Start with a defined future outcome and plan backward so that scope and investment follow evidence.

Worked example: Project Swift

The team planned backward from the one-year goal rather than forward from the prototype.

Outcome (about one year out). The 150-store regional deployment holds counter time at or below 3:30 and completes at least 80 percent of returns without a manager. Arrival signal: two consecutive months of regional dashboards at or above target.

Milestone (launch). The 50-store launch reaches activation at or above 80 percent in live stores by Month 13, with average counter time at or below 3:30. Evidence: the Tranche 2 release criterion, recorded as met (appendix A7). Prerequisite: the pilot milestone, which released Tranche 2 at Month 10.

Milestone (pilot). The guided-returns path clears the same thresholds in a 10-store pilot. Evidence: pilot metrics at 3:18 and 81 percent. Prerequisite: the prototype milestone, and a chartered sandbox for the 10 stores (appendix A10).

Milestone (prototype). A prototype confirms the need with target users. Evidence: at least 12 of 20 target users confirm a critical need. Prerequisite: the Wizard of Oz trial confirms the time saving.

Smallest next step (now). Run the early guided-returns trial as a Wizard of Oz test at a single store, with a person performing the policy check behind the guided screen, to confirm the time saving before any build (chapters 2 and 6).