

OPERATING CASE / OC01

Plant Operating Excellence

Select improvements by finished output and economic alternatives

A Luna Sol operating design for a plant improvement mandate: diagnose the constraint, choose a bounded Kaizen portfolio, build frontline capability and verify customer-level value.

PLANT TRANSFORMATION / ILLUSTRATIVE OPERATING SCENARIO

BLUF / BOTTOM LINE UP FRONT

Improve the constraint.
Protect the customer commitment.

MODEL-DERIVED DECISION

The modeled preparation-speed project halves local hours but adds no good units. The quality option completes demand; its economic advantage depends on the strongest feasible alternative.

PROPOSED ENGAGEMENT RESPONSIBILITY

Proposed Luna Sol role: operational diagnosis, project selection, governance, capability transfer and evaluation.

EVIDENCE STATUS

Invented scenario and figures; reproducible analysis. No completed client engagement or realized benefit claimed.

IS: illustrative scenario • MR: model result • PD: proposed design • RB: primary research. Figures are modeled, not client results.

Choose the improvement that reaches the customer

Context. A plant improvement leader must select among competing Kaizen projects while preserving product commitments, quality and the availability of a small specialist team. Faster tasks and more completed workshops can create a convincing progress report while the limiting process and customer shortages remain unchanged.

Objective. Decide which action to test first, what evidence would justify scaling it, and who can authorize changes. The worked campaign contains two product families, preparation followed by finishing, explicit rework and a three-day improvement-team budget. These are invented operating conditions used to demonstrate Luna Sol's approach.

Approach. Map original demand through first-pass finishing, rework and final release. Enumerate every permitted batch mix and action combination, preserve the service floors, and compare contribution after incremental project charges. Link the planning decision to a field measurement contract, operator capability, a prospective evaluation and an evidence-gated handoff.

Findings. The preparation-speed option halves the preparation hours used by the selected plan, from 22 to 11, but leaves good output at 81 of 94 demanded units. Its project charge reduces contribution from 988 to 958 synthetic monetary units. The quality option produces all 94 demanded units and contribution of 1,116 under the declared model. A finishing-setup alternative yields 90 units and contribution of 1,018.

Decision value. The plant sponsor should prioritize validation of the quality mechanism and retain finishing setup as the credible alternative. With partial-service floors, the quality option loses its economic advantage when its campaign charge exceeds 178. The more permissive 208 break-even against doing nothing is the wrong investment threshold when the finishing alternative is available.

Implications. Improvement selection belongs at the complete-flow level. Production owns the feasible release plan; Quality approves the process change; Finance validates avoidable economics; the site sponsor resolves shared capacity. Expansion requires a measured mechanism, preserved service and demonstrated unaided execution.

Classification: illustrative scenario (IS) = invented context and assumptions; model result (MR) = recalculated from those assumptions; proposed design (PD) = future delivery and evaluation; research basis (RB) = cited primary methods. This is an operating-design case, not a record of a completed client engagement.

Exhibit 1 | Original demand defines the plant problem

[IS] The campaign has 54 demanded units of A and 40 of B. The main comparison protects minimum completions of 45 A and 36 B; the unfilled balance remains visible. A second analysis requires all demand. Neither condition may be silently substituted for the other.

Per ten-unit batch	Product A	Product B
Preparation hours	2	3
Primary finishing hours	3	4
First-pass good / rework attempts	8 / 2	7 / 3
Recovered good / scrap after rework	1 / 1	2 / 1
Hours per rework attempt	0.5	0.5
Material cost per started unit	8	12
Revenue per demanded good unit	18	30
Variable cost per rework attempt	1	1

[IS] Preparation has an earlier booked 24-hour window; finishing has a later booked 42-hour window. Every selected preparation batch can be completed before finishing begins. Improvement work has three team-days available. There is no opening inventory, overtime, outsourcing, carry-in or incomplete batch credited as finished output.

[MR] Current-method A consumes four finishing hours per batch after rework; B consumes 5.5. At five A and four B batches, finishing consumes 42 hours. Preparation uses only 22. The two unused preparation hours do not solve the finishing constraint.

Source: frozen authored scenario, inputs.json. All monetary figures use SYN, a synthetic accounting unit. They are neither quoted client prices nor reported profit. Fixed booked-resource payroll and overhead are common and excluded.

Exhibit 2 | Local speed can reduce contribution

[MR] Enumeration evaluates all eight action subsets and 35 demand-bounded batch mixes: 280 candidates. Eight satisfy every team, resource, service and demand constraint. The table gives the highest-contribution plan for each feasible subset. Batch counts are A/B; good units are A/B; hours are preparation/finishing.

Action	Batches	Good units	Hours	Net SYN
Current method	5 / 4	45 / 36	22 / 42	988
P: preparation speed	5 / 4	45 / 36	11 / 42	958
F: finishing setup	6 / 4	54 / 36	24 / 42	1,018
Q: B quality	6 / 4	54 / 40	24 / 42	1,116
P + F	6 / 4	54 / 36	12 / 42	988

[IS] P halves preparation time for a charge of 30 and one team-day. F removes one primary finishing hour per B batch for a charge of 50 and two days. Q changes each B batch to nine first-pass good units plus one recovered unit, with one rework attempt, for a charge of 80 and three days. The model composes changes without extra synergy; combinations containing Q exceed the three-day budget.

[MR] P delivers the same 81 good units as the current method while adding its 30 charge. Combining P with F uses the whole team budget but adds no good units over F. The additional 12 preparation hours released by P+F are not counted as payroll savings, saleable output or cash.

[MR] Q and F use the same six-A/four-B start plan. Q produces four additional B units and uses eight fewer rework attempts than F. At the declared prices and charges, that difference is $4 \times 30 + 8 \times 1 - (80 - 50) = 98$ SYN. This explains the choice without attributing every difference to generic productivity.

Source: calculate.py/results.json; independent Fraction-based reconstruction uses a wider 968-candidate grid and recovers the same eight feasible assignments. Model feasibility is not field validation.

Exhibit 3 | The alternative sets the economic hurdle

[MR] Quality produces pre-charge contribution of 1,196; finishing setup produces 1,068. With F’s charge fixed at 50, Q is preferable only while 1,196 – cost(Q) exceeds 1,018. The switching charge is 178. Comparing Q only against current-method contribution of 988 would permit a charge of 208 and approve some inferior choices.

B unit revenue	Q charge 80	Q charge 178	Q charge 220
24 SYN	Q: 876 vs F: 802	F: 802 vs Q: 778	F: 802 vs Q: 736
30 SYN	Q: 1,116 vs F: 1,018	Tie: 1,018	F: 1,018 vs Q: 976
36 SYN	Q: 1,356 vs F: 1,234	Q: 1,258 vs F: 1,234	F: 1,234 vs Q: 1,216

[MR] The charge threshold changes to 154 at B revenue of 24 and 202 at B revenue of 36. These are deliberate scenario changes, not a confidence interval or a sensitivity to observed market prices. Material costs, recipes, resource windows and demand remain fixed.

[MR] Full service changes the feasible set. With all 54 A and 40 B required, F cannot qualify because its B lots yield only nine good units and five lots would exceed the 40-unit sales ceiling. Q is the only feasible option under the three-day budget. A higher Q charge therefore triggers a contract, schedule or resource decision; it does not make F a compliant substitute.

[PD] The sponsor must first sign the service contract and then select the best feasible economic alternative. If full demand is non-negotiable, test material supply, rework validation and resource feasibility before committing the release plan. Do not describe a profitable partial-service plan as fulfillment of the whole customer obligation.

Source: independently reproduced cost/price sensitivity. Main-case floors permit unmet demand; complete-service sensitivity does not. Revenue is assumed collectible in the planning model; no actual sale or cash receipt exists.

Exhibit 4 | Diagnose the constraint at the work

[Proposed illustrative design] The plant sponsor’s first decision is which loss mechanism deserves a controlled trial. The example follows a mixed-product flow from preparation into constrained finishing, quality disposition and release. A faster preparation step is useful only if the rest of the flow can convert that work into conforming output when customers need it. No real plant observations or achieved improvement are asserted here.

Start at the work, with an operator and the relevant production, quality and maintenance owners. Trace selected original orders through both normal and disrupted conditions, across product families and shifts. Capture material availability, setup, active processing, blocked time, equipment stops, inspection holds and rework. Keep the original order, lot, product, quantity and due time attached to every observation; use split and merge links when a lot changes shape.

Observation	Evidence to collect	Decision owner and gate
Preparation finishes; work waits	Queue-entry time, finishing readiness and downstream disposition	Production lead: test the downstream constraint before increasing release
Finishing repeats or rejects work	First-attempt result, defect category, recipe version and disposition	Quality lead: authorize the cause test and any process-setting change
Changeover exceeds the plan	Last conforming piece, readiness checks and first conforming piece of the next run	Cell lead: separate missing prerequisites from the physical setup task

Use a defined primary time state so the same minute is not charged twice as both a stop and a quality hold. Preserve secondary cause flags when they coexist. Reconcile equipment timestamps with a bounded manual observation sample; a blank event field means unknown, not zero loss. Retain the observer’s evidence and the codebook version when supervisors disagree about classification.

The diagnostic handoff is a current-state flow map, a reconciled loss ledger and a short cause register. Each proposed cause states what would be observed if it were correct and what evidence would contradict it. For example, if finishing starvation is the claimed constraint, observations of finishing being continuously busy while prepared work accumulates challenge that explanation. A workshop vote or a dashboard correlation does not resolve the fork.

The sponsor approves a trial only when its affected work population, owner, prerequisite evidence and stop route are clear. Maintain safe existing work while information is gathered; no improvement workshop authorizes bypassing established quality release or equipment safeguards.

Source basis: the supplied role brief establishes the requested diagnostic capability. This page is a proposed plant operating method; the separate numerical scenario uses invented inputs and does not supply shop-floor evidence.

Exhibit 5 | Give each improvement a bounded contract

[Proposed illustrative design] The portfolio gate compares preparation speed, finishing setup and first-pass quality improvements under a limited improvement-team budget. Do not rank them only by local minutes saved. Each charter must explain how its changed work reaches conforming customer output, which scarce specialist it consumes, and which product-specific service requirement it must preserve. A project that expands preparation capacity may have little near-term value when finishing remains constrained.

Candidate package	Operating hypothesis	Release evidence
Preparation readiness and setup	Stage approved materials and tools before the changeover; reduce avoidable waiting	Matched setup definition, readiness checklist and no transfer of unfinished work to another shift
Finishing setup	Reduce finishing setup time without changing product acceptance	Matched setup boundaries, equipment readiness and authorized method
First-pass quality	Reduce repeat processing within the approved process window	Quality-approved trial, reconciled rework and unchanged acceptance standard
Daily flow and release discipline	Match work release to usable finishing capacity and the original order priorities	Visible queue, product-family service floors and a named exception decision

The operating model gives the plant manager authority over priorities and protected customer commitments. The process owner owns standard work; Quality owns acceptance and disposition; Maintenance owns equipment changes; the planning lead owns the demand and sequence basis. The improvement lead facilitates analysis and coordinates delivery. Finance validates the benefit category and cost basis rather than inheriting the project team’s largest estimate.

Before scheduling events, reconcile the same engineer, trainer and operator release time across every charter. Planned attendance is not available capacity when the same person is assigned to concurrent trials. Sequence dependencies explicitly: measurement repair may precede a process experiment, and equipment readiness may precede new standard work. A project can be paused for a missing prerequisite without being declared economically worthless.

At the proposed weekly portfolio gate, review one evidence packet per active project: original scope, current hypothesis, completed test, deviations, remaining resource demand and the next release decision. Use continue, revise, hold or close dispositions with named ownership. Close an administrative action when its evidence is received; close a process issue only when the agreed operating acceptance test is met.

Preserve the current method and low-effort alternatives as genuine comparators. If preparation speeds up while downstream inventory grows and mature good output does not improve, report that pattern directly. It may reveal a constraint or a sequencing need; it is not verified enterprise productivity. The numerical scenario tests declared capacity and yield assumptions, while field selection still requires observation, safe experimentation and local authority.

Source basis: proposed Kaizen governance for the illustrative preparation–finishing flow. Project selection, staffing availability, improvement effects and financial realization have not been observed in a client plant.

Transfer a working routine, not just a completed training deck

[Proposed illustrative design] Sustainable handoff requires the receiving team to perform, recognize an exception and recover within its authority. The capability plan therefore pairs each changed task with an observable demonstration. Attendance is useful administration but does not prove competent execution. The sponsor should release the next rollout stage only after the process owner accepts both the work method and the conditions under which it must stop.

Receiving role	Demonstration in a safe rehearsal	Acceptance evidence
Operator	Perform the approved preparation or finishing sequence and explain when to hold	Observed task record linked to the current instruction and equipment version
Shift leader	Detect a missing readiness record, preserve the order and route the exception	Correct hold, owner assignment and escalation without altering the due time
Process owner / supervisor	Review a deviation and restore the authorized method	Disposition, document change record and follow-up observation
Plant sponsor	Resolve conflicting project and production resource demands	An explicit priority decision with service and workforce consequences

Build training around the work that actually changes: tool and material preparation, first-piece checks, defect escalation, rework routing and queue handoff. Use supervised practice or a simulation appropriate to the risk. A deliberately missing release record can be rehearsed without producing unsafe material. Include relief staff and the shifts that will inherit the method; a successful demonstration by the project team alone is not plant capability.

Use the proposed shift review to reconnect the standard with reality. Operators surface deviations; supervisors distinguish a missing resource from a misunderstood instruction; the process owner decides whether the method needs revision. Retain the issue, assigned owner, interim containment and closure evidence in a versioned log. Staff should be able to raise a problem without being pushed to relabel it merely to protect a project’s reported benefit.

The handoff packet contains approved standard work, applicability by product and equipment, the measurement dictionary, demonstrated role coverage, spare or fallback requirements, and open exceptions. It also identifies who can change settings, approve product disposition and authorize restart. A changed material, recipe, tooling arrangement or staffing assumption triggers review of the affected controls; it does not automatically invalidate unrelated work.

Taper coaching only when the receiving leaders can run the routine and produce the required evidence without the facilitator completing it for them. Observe follow-up under normal workload, not just a staged event. Share the method across plants with its prerequisites and failed conditions attached. Transferable learning is a testable mechanism; copying a checklist or a success label is not proof another site can reproduce the result.

Source basis: proposed coaching, adoption and handoff design. No completed training, operator certification, plant rollout or personal professional credential is asserted.

Define the metric before claiming the benefit

[Proposed measurement design] Freeze definitions before a trial: observation unit, original population, event time, product family, process version and owner. Completing a Kaizen event, following an instruction and producing acceptable output are different achievements; none can substitute for the others.

Measure	Definition and preserved basis	Owner / decision use
Good finishing throughput	Accepted finishing units / scheduled production hours in a fixed window; report each product family	Production: locate the constraint; do not sum incompatible units
First-pass finishing yield	First-attempt passes / first-attempt units at finishing; report retests separately	Quality: preserve acceptance criteria
Rework and scrap	Repeat-processing work and scrapped original units; retain stage, identity and final disposition	Quality: retain upstream losses and hidden repeat work
Original due-unit service	Conforming quantity by original due time / originally due quantity; cap credit at each order's requirement	Planning: preserve family floors; excess cannot offset missed work
Constraint loss time	Exclusive primary elapsed states in the same scheduled window; retain secondary cause flags	Cell lead: reconcile shortage, blockage, stop and hold
Economic realization	Separate released capacity, avoidable cost and cash; include trial and sustained-control work	Finance: approve value basis; prevent double counting

Reconcile original quantities into accepted, scrapped and still-open states at each cutoff, retaining the rework history. Do not erase held material or unresolved orders when a shift ends. A unit accepted after its original due time improves eventual completion, not its earlier service result. Use a later, declared maturity check to resolve final disposition; keep the interim report explicitly provisional.


For comparisons, preserve raw counts and exposure as well as rates. Stratify by product family and relevant process version before combining results. If mix changes, show both actual-mix performance and a clearly defined fixed-mix sensitivity. A better blended average can arise from different work rather than a better process. Missing outcomes need an explicit unknown category and follow-up owner; excluding them cannot manufacture a favorable result.

The release gate requires a reproducible calculation, reconciled records and the agreed service, quality and operational safeguards. A short favorable run is not proof of sustained control or causality. Statistical evaluation and economic break-even belong to the separate analytical design. This dictionary makes those tests interpretable; it does not supply the missing baseline, validate a sensor or certify the proposed change.

Source basis: proposed field measurement contract, broader than the simplified campaign model. Approve definitions against local systems before use; modeled service covers declared campaign windows, not a detailed customer-deadline simulation.

Exhibit 6 | Separate exposure, mechanism and outcome

[PD] The diagram is a causal hypothesis for a future field study. Exposure to coaching is distinct from implementing the frozen action package. The worked model validates its own arithmetic; it does not identify the causal arrows below.



Assignment → exposure → action. Record which cluster was offered which package, actual support hours, the version used and any crossover. A treatment-assignment effect includes imperfect adoption. A per-protocol association among enthusiastic adopters is a different quantity and may reflect selection.

Action → mechanism → outcome. Link rework attempts to original units, total finishing hours and accepted output at the campaign cutoff. The corresponding customer outcome is completion against the original obligation; in a real plant, subsequent complaints and returns require a further declared maturity window.

Context and omitted paths. Product difficulty, skill coverage, maintenance, incoming material quality and concurrent capital work may affect treatment choice and results outside random assignment. Record shared supervisor and machine interference. Residual differences could arise through unmeasured paths; the diagram does not establish which explanation is responsible.

[PD] The primary comparison for investment is assignment to Q versus F, the strongest feasible alternative under the main service contract. A Q-versus-current-method comparison answers the separate total-program question. Estimate total assignment effects without controlling for post-treatment rework or adherence.

Detect the economic threshold before promising an answer

[PD] For comparable eligible line/crew clusters, define the primary estimand as the mean Q-minus-F difference in pre-charge campaign contribution over a prespecified evaluation window. Fix original demand, recognized good units, variable inputs and maturity rules. The net investment contrast subtracts the incremental Q campaign charge. Under the model, the gross contrast is 128 SYN, the incremental charge is 30, and the net contrast is 98.

[PD] Pair comparable independent lines or crews and randomize Q or F within pairs, after engineering and quality eligibility review. Collect repeated baseline campaigns, then aggregate each cluster over a fixed number of evaluation campaigns. Analyze independent pair-level contrasts. Repeated shifts on one line do not become independent clusters; shared staff, tools and learning may require a larger assignment unit.

Assumed SD of pair contrasts	Detectable margin above hurdle, 12 pairs	Pairs for a 30-SYN margin
45 SYN	36.40 SYN	18
80 SYN	64.71 SYN	56
120 SYN	97.06 SYN	126

[MR/PD] The investment rule requires the lower endpoint of a two-sided 95% confidence interval on gross effect to exceed 30 SYN. Screen 80% power at a design alternative of 60: the margin above the hurdle is 30. Required pairs = ceiling[(((1.96 + 0.842) × SD / 30)²]. With 12 pairs and SD 80, the detectable margin is 64.71, so gross effect would need to be about 94.71. Power to distinguish an effect from zero does not establish that it clears the economic hurdle.

[PD] Redesign before enrollment: recruit more independent eligible clusters, reduce measurement error, or use prespecified baseline adjustment supported by pilot covariance. Estimate serial and within-cluster covariance, simulate the actual allocation and follow-up, and use finite-sample critical values. The counts above are a normal screen, not a completed power analysis or evidence that the required network exists.

If a feasible redesign cannot resolve the investment decision at the chosen alternative, limit the pilot to mechanism and delivery feasibility. Preserve uncertainty and use a reversible staged decision; do not interpret non-significance as zero benefit. Planning arithmetic: economics.py.

Exhibit 7 | Buy information only when it changes the choice

[IS] A pre-campaign diagnostic asks whether Q will work. Success returns 1,116. If mandatory qualification detects failure before production release, the original recipe and a five-A/four-B fallback return 988 less the sunk 80 charge, or 908. Q’s charge includes that qualification; its three team-days are spent. This fallback preserves partial-service floors only. F returns 1,018 in both states. Diagnostic sensitivity and specificity are each assumed 80%.

Prior probability Q works	Best action before test	EVPI	EVSI	Net EVSI at fee 10
20%	F	19.60	0.00	−10.00
50%	F	49.00	28.20	18.20
80%	Q	22.00	1.92	−8.08

[MR] At the 50% prior, a positive signal moves the probability to 80% and selects Q; a negative signal moves it to 20% and selects F. Both signals have probability 50%. Expected value after the signal is 1,046.20, versus 1,018 without it: EVSI is 28.20. Fees of 30 and 50 turn net EVSI into −1.80 and −21.80.

[MR] At the 20% prior, even a positive signal only reaches a 50% probability, below the Q choice threshold of 110/208, or 52.88%. The same nominally accurate diagnostic therefore has zero decision value. Test accuracy and information value answer different questions.

[PD] The diagnostic occurs before option selection and uses resources outside the campaign team. Failed Q qualification must be known before committing the successful six-A/four-B mix, which is infeasible under the failed recipe. If testing consumes team-days, qualification is late, or full service is mandatory, rebuild the actions and payoffs. A fee of 30 is rejected at the three displayed priors, not at every possible prior.

EVPI is the value of perfect state information; EVSI is the expected improvement in the decision policy under the declared signal. Checks enforce $0 \leq EVSI \leq EVPI$. Do not add either quantity to Q’s 128-SYN comparison with the current method or its 98-SYN advantage over F.

Improve the active constraint, then test the whole campaign

[Proposed design] The Constraint-to-Capability Flow Contract (CCFC) links a plant improvement to original demand, scarce processing work and an executable operator method. Its first test is system-level: can the proposed change improve an authorized campaign outcome without violating product-level service, quality or improvement-team limits? A faster local task is a mechanism candidate, not a finished-goods result. MIT's capacity lecture separates process mapping, capacity, congestion and financial analysis; these questions should not collapse into one efficiency score [R1].

[Analytical inference] In the declared two-product campaign, preparation has an earlier booked window and finishing a later one. All required preparation can precede finishing. The model therefore tests resource feasibility rather than a detailed shop-floor queue, overlapping shifts or customer-delivery schedule. A preparation-time reduction expands one resource constraint; it need not improve the best feasible contribution when another constraint still binds. Any released preparation hours remain capacity unless a separate, credible cost-avoidance mechanism exists.

Decision object	Required interpretation
Good output	Count finished acceptable units against each product's original demand; preserve unmet units rather than removing them from the denominator.
Scarce finishing work	Include primary processing and every rework attempt, including unsuccessful attempts. Recovered units are not additional original demand.
Campaign contribution	Apply the declared sales and variable-cost assumptions once. Keep common paid resource windows and hypothetical option charges distinct from cash released.
Operator capability	Observe correct task execution and escalation after coaching tapers; a signed training record establishes neither sustained performance nor a credential.

[Proposed design] Freeze the action package before testing it: approved method, prerequisites, resources, quality checks, authority and response when the process deviates. NIST's Job Instruction sequence provides a useful preparation, demonstration, trial and follow-up structure [R2]. This supports a testable training protocol, not a claim that this case has trained plant employees.

[Boundary] Separate product-specific first-pass yield, recovery, scrap and finishing hours. Improving first-pass quality may both reduce finishing load and change saleable output; charging the same avoided rework again as a separate throughput benefit would double-count it. The constructed comparison identifies a result only within its declared deterministic recipes, demand ceilings and feasible choices. It establishes no realized client benefit or transferable yield forecast.

Evaluate the improvement package on original demand and independent crews

[Proposed design] The field estimand is the effect of assignment to a frozen improvement package, compared with the strongest feasible alternative (Q versus F for the primary investment test, with current method as a separately defined total-program comparator), on good units completed within an agreed campaign and on net avoidable campaign cost. Fix the production-order population, product specifications, resource calendar and completion boundary before assignment. Keep rejected, scrapped, reworked, unfinished and unmet units visible. Report product-level service beside the aggregate; improved total output can conceal a breach for one product.

[Proposed design] Use comparable line or crew clusters where training, shared tools and supervisors create spillovers. Match or stratify on preassignment product mix, equipment configuration, skill coverage and baseline performance, then randomize the package when operationally feasible. Do not randomize individual pieces while a common crew applies the new method to all of them. If only one line is available, a proposed phased comparison must address learning carryover, maintenance and demand changes; a simple before/after result remains descriptive.

Measure	Validation and decision use
Accepted unit and rework status	Freeze specifications and routing rules. Reconcile original IDs, attempts, recovered units and scrap; independently adjudicate ambiguous states.
Time and dimensional quality	Check clock boundaries and instrument repeatability, operator differences and stability. Do not attribute a gauge or logging change to the process [R4].
Capability transfer	Assess an authorized task without live coaching, then follow performance after support tapers. Record assistance and escalation; attendance alone is insufficient.
Economic effect	Value only the identified outcome and avoidable resource consequences. Include training, changeover and recurring support; service and safety gates remain binding.

[Proposed design] Analyze assignment effects at the cluster level. Realized adherence, training completion and rework are potential mediators; do not condition on them to claim a total treatment effect. Preassignment mix and configured equipment can be moderators. Preserve shared-resource interference and concurrent maintenance or capital changes in the record. Compare sustained unaided performance, not only the coach-supported launch period.

[Precision boundary] A field test needs baseline cluster variance, serial dependence, feasible recruitment, an economic hurdle and an acceptable service margin. A separate model-specific planning calculation can illustrate detectability under assumptions; invented campaign coefficients cannot estimate field power. A confidence bound that does not clear the hurdle is a decision hold, not proof that the true benefit is zero. Engineering and worker-safety approval remain separate from statistical precision.

Separate process signals, selection effects and capability claims

[Proposed design] No observed chronological plant panel accompanies the opportunity brief. Integer campaign alternatives, product rows and sensitivity points are not consecutive process observations, so an XmR chart should not be manufactured from them. For a future commensurate chronological series, document spacing, product mix, phase, rational subgroup and serial dependence before charting. Check measurement-system changes first [R3, R4].

[Recalculated method] With suitable individual observations, define $MR(t)=abs[x(t)-x(t-1)]$ and tentative individuals limits as $mean(x)$ plus/minus 2.66 times $mean(MR)$; moving-range limits are zero and 3.267 times $mean(MR)$. At least two observations are needed even to compute a moving range, while sparse, trending or definition-changing data make the fitted limits non-diagnostic. A point inside those limits does not establish improvement, capability to a specification or intervention causality. A rate chart requires authentic numerator and denominator counts; rounded yields and deterministic conversions do not create binomial data.

[Analytical inference] Selecting the worst recent line creates regression-to-the-mean risk. Under an explicitly assumed linear conditional-mean model, expected untreated follow-up is $m + \rho(x-m)$, where x is the selected baseline, m the comparison-population mean and ρ persistence. For an adverse high baseline, expected untreated decline is $(1-\rho)(x-m)$. Evaluate a plausible reliability range if data are absent; do not label its assumed reversion share estimated. Prefer concurrent eligible comparisons and repeated preperiod observations.

Falsifiable proposition	Disconfirming observation and scope
P1: The proposed lever improves the constrained campaign	At the same original demand, resources and service floors, a separately reproduced feasible comparison provides no claimed system gain. Local speed or utilization alone cannot rescue the claim.
P2: The quality change improves the finishing conversion	Linked records show that total finishing work per accepted good unit or product service fails its prespecified improvement threshold under valid measurement. A precise adverse result refutes; wide uncertainty warrants a hold.
P3: Coached capability persists without the coach	Post-coaching task assessments or mature quality/service results fail the prespecified retention criterion in comparable work. Attendance and successful coached trials cannot substitute.

[Boundary] CCFC is a proposed framework. The opportunity specifies a role need; it does not establish Ryan's credentials, selection, plant deployment or causal results. A reproducible illustrative planning result can demonstrate reasoning while implementation, workforce capability and sustained customer value remain unestablished.

Validate classifications and preserve uncertainty

[PD] First-pass quality, rework and scrap need consistent classification, not only a clean dashboard. Before rollout, draw a stratified random audit of 120 original units across product family, shift and proposed status, oversampling rare scrap and disputed rework. Two trained raters independently classify de-identified records using a frozen codebook, blind to the assigned package and each other's result. Retain the probability of selection so an oversampled audit is not reported as plant prevalence.

[PD] Report the complete confusion matrix, class prevalence, raw agreement, Cohen's kappa with an appropriate uncertainty interval, and unweighted versus design-weighted summaries where relevant. Adjudicate after independent coding and record changes to the codebook. Kappa is sensitive to prevalence; agreement does not prove engineering validity. A 120-record pilot is a proposed starting allocation, not a justified precision guarantee. Use pilot disagreement and clustering to size the definitive audit.

[IS/MR] A separate selection sensitivity illustrates regression to the mean. Suppose an adverse loss index is 100 for a selected line and 80 in its comparison population. Assumed persistence values of 0.25, 0.50 and 0.75 imply untreated follow-up of 85, 90 and 95 under $m + \rho(x - m)$. A decline of 15, 10 or 5 points could therefore occur without treatment under those assumptions. These are hypothetical conditional means, not observations or an SPC series.

Identification and measurement limits. There is no observed plant panel, deployment, randomly assigned intervention or verified customer outcome. Deterministic batch recipes omit stochastic yield, detailed sequencing, setup interactions, maintenance and incoming-material variation. Booked sequential windows establish a planning boundary, not a validated delivery schedule. Real specifications, instrumentation and coding require independent validation.

Transportability, attribution and economics. Two products and a three-day team limit cannot establish an industry benchmark. Model charges are campaign-specific; common payroll is not avoided. Revenue assumes demanded good units can be sold. No price, probability, covariance or value is estimated from a client. Luna Sol's proposed leadership role is design, coordination and decision discipline; this document establishes no historical plant engagement, certification or field outcome.

Numerical privacy: every scenario number is authored, so no confidential client metric or baseline is being reconstructed. Original opportunity titles, source locations and qualification requirements stay in the private evidence pack.

Primary research and the case contribution

[Research basis] Primary sources inform the methods and operating design. They do not establish the synthetic findings, certify the local implementation or supply client outcomes.

R1 | Class 3: Capacity Lecture, 15.760B Introduction to Operations Management [Class 3: Capacity Lecture, 15.760B Introduction to Operations Management](#)

Assess preparation and finishing jointly; separate scarce finishing hours, original demand and campaign contribution. A local speed gain is not automatically a system throughput gain. Teaching methods do not validate this invented plant or its coefficients. OC01 uses sequential booked resource windows, not a stochastic queue or empirical Little's-law estimate. None of the lecture's example quantities or outcome claims are transferred.

R2 | Training Within Industry (TWI) [Training Within Industry \(TWI\)](#)

Separate attendance and documented standard work from demonstrated independent task execution. Define maintenance of the new method after coaching is reduced as an untested mechanism. Do not transfer the page's reported training-time, scrap, safety or cost benefits into OC01, promise the same effects, or infer Ryan has TWI/Lean/Six Sigma credentials.

R3 | e-Handbook: Individuals Control Charts and Assessing Process Stability [e-Handbook: Individuals Control Charts and Assessing Process Stability](#)

Assess chronological comparability, product mix, measurement changes and phase before tentative XmR. A finite authored scenario table is not a plant time series. Stable variation and conformance to engineering specifications are different questions. Do not treat a few points inside fitted limits as proof of stability or improvement; control limits are not service targets or causal confidence intervals. The source's sample-size example is not adopted as a universal minimum. No actual OC01 SPC panel currently exists.

R4 | e-Handbook: Gauge R & R studies; Analysis of repeatability; Analysis of reproducibility [e-Handbook: Gauge R & R studies; Analysis of repeatability; Analysis of reproducibility](#)

Before estimating quality or cycle-time improvement, verify the measurement process across instruments, operators, product families and shifts. Use separate independent classification checks for scrap/rework reason codes. Gauge variance formulas apply to suitable quantitative repeated measurements, not automatically to categorical reason-code agreement. No gauge study, instrument calibration, certification or field measurement validation has been performed for OC01.

[Contribution] This case does not claim a new theorem or a field extension of the literature. Its contribution is a reproducible decision example: the strongest alternative tightens the quality-project hurdle from 208 to 178, and a full-service contract can remove that alternative entirely. CCFC combines system-flow accounting, explicit decision rights and a test of unaided capability; its propositions remain open to field falsification.

Internal referee and partner review

Review scope. Author-led review with separate agent roles for primary-source methods, operating design and numerical analysis; the coordinating author independently rebuilt the finite model using exact fractions. This is not external peer review, client validation or a certification of realized impact.

Major issue	Resolution and surviving boundary
Local metric presented as system value	Report original demand, finishing load and net contribution together. P releases preparation hours but adds no finished units in the finite model.
Weak comparator inflates willingness to pay	Show 178 versus F and 208 versus the current method. Complete-service feasibility is a separate decision contract.
Resource totals mistaken for a schedule	Disclose sequential booked preparation and finishing windows. Detailed dispatching, stochastic queues and actual deadlines remain outside the model.
Small pilot claimed to prove economics	Publish the variance scenarios and MDE/hurdle mismatch. Require design-specific simulation and feasible cluster recruitment before a confirmatory study.
Illustrative material presented as experience	Declare invented figures and proposed operating responsibilities. No client, implementation, credential or realized saving is claimed.

Minor issues addressed. Preserve original quantities across rework, avoid adding information value to operating benefit, distinguish attendance from independent capability, and give the finance, quality and production owners explicit release gates. Reference methods support the approach; their institutional outcomes are not imported.

Partner critique. Page one supports a bounded decision: validate the quality lever first, keep finishing setup as the alternative, and reject a preparation-only project as a throughput investment under the declared conditions. The COO must approve the service obligation before using the economic threshold. A real commitment still requires a verified recipe, resource plan and quality authority.

Disposition. Suitable as an illustrative operating-design case with auditable planning calculations and proposed delivery mechanics. Historical implementation and causal outcome impact remain unestablished. The private annex provides the assumption ledger, executable model, independent reconstruction and review record.