



Elganainy Adaptive Marketing Framework

EAMF: From understanding individual needs to shaping value around them

A proposed applied framework attributed to Karim Elganainy. Working version, 16 October 2026. The name has not been commercially or legally cleared. This is not a validated scientific theory or a claim to have invented individualization. All cases and numerical examples are hypothetical. The stages constitute a proposed synthesis requiring testing; the cited sources do not validate the framework as a whole.

Proposition and scope

AI may reduce the cost of understanding a need and configuring some products and services. A useful marketing question becomes: "What feasible configuration helps this person achieve their intended outcome in this context?" This does not mean previous marketing always sold one product to everyone. Customization and made-to-order services already exist. The proposition concerns the potential speed and reach of adaptation, connected to product delivery and measurement rather than advertising alone.

The decision unit is person $\times$ context $\times$ moment $\times$ task, rather than a fixed identity that a company claims to understand completely. The same person can need different configurations. Use the least personalization, data and complexity necessary to create real value.

Begin with a good standard baseline, then distinguish five potentially overlapping layers:

1. Communication personalization: change the message or explanation.
2. Recommendation personalization: select or rank existing solutions.
3. Experience adaptation: change the interface, workflow or service experience.
4. Offer configuration: assemble a product or bundle from approved modules.
5. Core product generation or modification: create part of the functionality, content or configuration itself within testable limits.

These are not levels every organization should climb. Choose the simplest layer that serves the need. Digital, physical and clinical settings impose different constraints; they are not higher levels of personalization. A model's ability to describe a physical or clinical change does not authorize its execution.

Working Backwards from Need starts with an intended customer outcome and derives solution requirements. Reverse Engineering examines an existing product or system to understand its structure or operation. Reverse engineering can be a supporting tool, but is not an accurate name for the whole framework.

For customer-built or co-designed solutions, ask which role the customer wants: accept a finished result, configure it, extend it through a permitted integration, participate in guided creation, or maintain a licensed fork. Compare buying, configuring and building on total lifecycle cost, reliable task completion and the customer's preferred level of involvement. Assign ownership of testing, compatibility, rights review,



updates and support before delivery. Participation is optional; do not transfer unwanted work or risk to the customer.

The seven stages

1. Understand with Permission

What outcome is intended, and what do we actually know? Start with the person's stated task, time constraints and preferences. Distinguish explicit statements from observed behavior and AI inferences. Give each signal a source, purpose, confidence, expiry and permitted use. Ask when uncertainty changes the decision. Do not collect sensitive information speculatively. Permission to provide a service does not automatically authorize advertising or model training. Offer a good standard path to people who decline personalization.

Deliverable: a brief stating "I want... under these circumstances... and success means...". Gate: necessary, permitted and sufficient information; otherwise ask or use the standard path. Owners: customer research and data/privacy responsibility appropriate to risk.

2. Frame Value and Constraints

Define a customer outcome and a deliverable promise. Specify what may change and what stays fixed. Separate hard constraints from preferences: safety, baseline quality, rights, local requirements, stock, capacity, price and cost-to-serve. Stronger persuasion cannot fix an unsuitable product. Set customer-value and business-feasibility criteria without reducing customer welfare to profit.

Deliverable: an internal value contract recording the promise, adaptation limits, escalation and stopping rules. Gate: a real need that can be served at acceptable risk; otherwise do not personalize. Owners: product, operations, finance and the relevant domain specialist.

3. Compose Feasible Options

Build an approved module library: functions, content, human services, appointments, sizes and channels. Every module needs cost, capacity, incompatibilities and version information. A model can propose options; independent rules verify feasibility, availability and price. Generated text must not become operational ground truth. Show a small set of meaningfully different alternatives alongside a standard option. Open-ended generation is unnecessary when simple rules suffice.

Deliverable: configurations with identifiers, cost, delivery timing and explanations of fit. Gate: compatibility, safety and economic feasibility; failed configurations never reach the customer. Owners: product, operations and specialist reviewers.

4. Enable Choice and Control

Explain the recommendation, the assumptions behind it and what the customer can change. Show full prices and trade-offs. Provide rejection, reset and opt-out. Do not exploit fear or inferred vulnerability. Separate product selection from data permissions. A rejection may reveal a mistaken understanding of the need rather than a persuasion problem.



Deliverable: the person's choice, modifications and specific permissions. Gate: an understandable choice with a real alternative. Manipulative interfaces fail this gate even if conversion increases.

5. Deliver with Accountable Oversight

Link what was proposed, selected and delivered using consistent identifiers. Assign responsibility, escalation and a route back to standard service. Reviewers need time, information and authority to reject, rather than nominal approval duties. Test capacity pressure and record actual costs and exceptions. Do not promise human support that the team cannot provide.

Deliverable: a service delivered as promised and a sufficient, data-minimized exception record. Gate: quality, timeliness and safety; reduce or stop adaptation if these fail. Clinical decisions remain under clinical responsibility.

6. Experiment and Measure

Compare against a good standard offer and, where feasible, simple rules. Prespecify the hypothesis, primary metric, randomization unit, meaningful effect size, required observation period/information and guardrails. Measure customer outcomes, contribution after cost-to-serve, agency, trust and safety. Clicks and screen time are not intrinsic customer benefits.

Deliverable: an effect estimate with uncertainty and a decision to scale, modify or stop. Healthcare marketing experiments must not alter treatment or delay care. Clinical research requires its own appropriate pathway.

Test authorship separately from product quality. Compare a ready-made solution, a fully AI-generated solution, and guided customer co-design while matching usable outcome quality as closely as possible. Track completion and abandonment, successful and failed construction, perceived contribution, psychological ownership, satisfaction, willingness to pay, support burden, and total cost. Record whether customisation replaces a branded purchase or complements it. A positive result among successful builders alone is not enough: include people who quit, and distinguish participation benefits from better fit or novelty. Psychological ownership describes a feeling of authorship, not legal ownership or permission to use a trademark. Treat any participation benefit as a hypothesis to measure, not a guaranteed IKEA effect.

7. Learn, Refresh and Retire


Update explicit preferences without converting one choice into a permanent identity. Give inferences an expiry. Separate immediate service adjustments from reviewed model updates and strategic portfolio reviews. Monitor the feedback loop: displayed options shape clicks, so clicks are not independent evidence of every underlying preference. Diversify low-risk options when appropriate and preserve reset rights. Monitor drift, versions and rollback capability. Deleting a record does not necessarily remove its influence from a trained model without a tested mechanism.

Deliverable: time-limited knowledge, documented updates and improved or retired modules. Governance, privacy and safety span every stage rather than appearing only at the end.



The EAMF operating loop

Proposed applied framework by Karim Elganainy. Each release must pass explicit feasibility and safety checks.



Across every stage: privacy • safety • feasibility • operational capacity • the right to refuse

Original explanatory diagram. The seven-stage EAMF loop is a proposed applied framework, not an empirically validated theory.

Quick stage-gate checklist

1. Understanding: has the person reviewed the need brief, and have statements been separated from inferences?
2. Constraints: are the outcome and non-negotiable elements explicit?
3. Configuration: have module validity, prices and capacity been checked?
4. Choice: can the person edit and decline without unnecessary friction?
5. Delivery: is there an accountable owner, time for review and a rollback path?
6. Evaluation: were the comparator, metrics, sample plan and guardrails specified before examining results?
7. Learning: is it clear when to review, delete or stop using retained knowledge?

STP, the 4Ps and the customer journey

Segmentation remains useful for market sizing, shared needs, module design, feasibility and fairness. Add dynamic task and context information. Targeting selects use cases the organization can genuinely serve, rather than everyone it can track. Positioning preserves a consistent brand promise and baseline quality while the configuration of value adapts. An individual is not a perfectly known “segment of one.”

Product becomes a platform of feasible modules and configurations. Price should expose components and trade-offs; initially hold pricing policy fixed to isolate product effects, and do not use vulnerability for hidden differential pricing. Place adapts channel, timing and fulfillment within capacity, with non-digital alternatives. Promotion explains what actually changes without unproven claims. Different advertisements alone do not establish different products.



The journey can become a loop: discover need → select/configure → use → review and adjust → receive support or leave. Comparison, purchasing and brand loyalty do not disappear.

Consumer behavior: conditional hypotheses

Expectations of fit may increase, and some people may participate in shaping a solution rather than compare finished products. Search effort may fall, or choice fatigue may rise. Price comparisons may become harder, increasing the importance of transparency. Dependency and lock-in may grow.

Personalization does not necessarily diversify demand. If many people use similar agents, ranking systems or data, recommendations and demand could concentrate on a small set of options. Test the diversity of options shown and chosen, concentration and the ability to override recommendations; do not present concentration as inevitable. An agent selecting on someone's behalf introduces another layer of decision-making power requiring control and accountability.

Sector limits

Digital products can change modules, interfaces and content relatively quickly, but inference, review, latency, quality and intellectual-property costs are not zero. Physical products require bills of materials, compatibility, inventory, manufacturing, testing and returns management; text generation does not eliminate production time. Human services face real staffing, training and capacity costs.

Healthcare experience personalization—language, appointments and explanations—differs from diagnosis, dosage and treatment. The latter needs clinical evidence, oversight and applicable authorization, not a marketing decision. Do not infer disease from advertising behavior to sell a test.

WHO supports autonomy, safety and rights in health AI. FDA explains that clinical decision-support classification depends on the software's functions and specific criteria; calling something an “assistant” does not determine its regulatory position. These are governance references, not legal advice or a transfer of US rules to Egypt or other jurisdictions. [1, 2]

Case: a free image editor with Photoshop-like functions

This is a hypothetical design, not a claim about an existing product's features. A designer wants to prepare ten store images with uniform dimensions, background removal and manual review. Instead of building an entirely new editor for each person, the system assembles an existing-module workspace, task sequence, shortcuts and export preset. It may propose limited, tested automation with a preview and undo. The core image engine and quality standard remain consistent. Free to the user does not mean cost-free to the provider: assess transparent funding rather than covert data monetization.

Permitted signals are the selected task, selected images and stated preferences. Do not automatically scan the entire computer or upload images to another provider without authorization. Show a plan and preview before destructive edits. Measure correct output, completion time, rework, file compatibility, processing cost and undo failures. Compare a standard template, simple configuration rules and AI. If rules win, do not claim AI is necessary.

The marketing difference is that the promise describes completion of the task, while the interface, workflow or automation adapts. AI's possible contribution to composing or generating workflows does not make free software, plugins or interface customization new inventions.



Case: games and mods

Modding, map creation and customization predate generative AI. The proposed addition is that an adult player describes a short cooperative session, lighter puzzles and a preferred difficulty. The system assembles a level from permitted elements or proposes a quest within engine rules, then checks playability, performance, safety and compatibility. Do not promise an automatically excellent infinite world.

Player controls include preview, acceptance, a fixed-level alternative and disabling adaptation. Developer controls cover content scope, age ratings, rights, anti-cheat and preventing untrusted code execution. Multiplayer requires group compatibility and fair shared rules; one player's preference must not secretly change the rules for everyone. Exploiting addiction susceptibility or spending is not the goal. Playtime alone is not value.

Measure reported enjoyment, completion of the intended task, playable variety, errors and crashes, fairness, generation/review cost, complaints and fatigue. Compare against good fixed content or appropriate human-created mods, not an intentionally weak baseline. Minors require additional safeguarding and governance; this initial example concerns adults.

Case: a clinic service, not automated treatment

An adult with an existing appointment selects the preferred language, channel and appointment window. The marketing layer does not need a medical history to show a time slot. Modules include available appointments, optional reminders, information formats approved by the clinical team and human contact. The system does not freely generate treatment or preparation instructions, or recommend tests to increase revenue. Health questions follow existing clinical escalation protocols, without delaying help to finish a marketing conversation.

Following appropriate institutional review, a study could compare standard and improved coordination while keeping care and required communications equivalent. Measure completion of necessary visits, comprehension, staff workload, care delays, privacy and fairness. More visits alone do not establish better health. Treatment personalization is a separate clinical domain.


Hypothetical calculation: 700 completed appointments per 1,000 eligible bookings in control versus 750 in treatment, a five-percentage-point difference. If each incremental visit contributes 20 currency units after variable care costs, potential additional contribution is 1,000. Additional coordination costs of 700 and review costs of 400 produce an effect of -100. Consider clinic capacity and displacement of other patients; individual randomization may be insufficient when capacity is shared.



Economics and formulas

Evaluate incremental value before scaling

A useful decision model, not a promise of uplift. Include delivery costs and customer harm.



A higher conversion rate can still be a worse business outcome.

Original explanatory diagram. Compare incremental value after costs with a suitable control group.

Contribution = realized net revenue – product/module costs – fulfillment – variable support and human review – variable AI costs – return and compensation costs not already accounted for.

Do not subtract refunds twice if they are already reflected in net revenue. State whether acquisition costs are included or tracked separately. This is not net profit: fixed costs remain separate.

Incremental contribution per eligible customer = average contribution in the adaptive group – average contribution in control. Include people who did not buy in the denominator. Contribution per order is a different diagnostic and cannot substitute for this measure.

Total feasibility over a stated horizon = eligible volume × incremental contribution per eligible customer – additional fixed setup, governance and maintenance costs over the same horizon.

Break-even volume = additional fixed cost ÷ positive incremental contribution per eligible customer. If the increment is zero or negative, volume does not solve the problem under this model.

Hypothetical adult-training platform: per 1,000 eligible people, control sells 100 orders with net revenue of 40 and service cost of 10 each, producing contribution of 3,000, or 3 per eligible person. Adaptation sells 130 orders with net revenue of 40 and service cost of 15, producing 3,250, or 3.25 per eligible person. The increment is 0.25. Fixed costs of 5,000 require 20,000 eligible people to break even if assumptions remain stable. If variable service cost rises to 18, contribution falls to 2,860: higher conversion, worse economics. These figures are illustrative, not observed results or evidence of statistical significance.

Balanced scorecard

- Customer: task success, time-to-value, output quality and preference for adaptation after use.
- Economics: contribution per eligible customer, cost-to-serve, returns, oversight costs and retention where relevant.



- Operations: valid configurations/all proposed configurations, fulfillment of promises, latency, escalations and reviewer changes.
- Agency: explanation comprehension, ease of editing and opting out, selection of standard options and pressure complaints. Few edits can indicate a difficult interface rather than satisfaction.
- Safety and privacy: confirmed incidents, out-of-scope outputs and data use outside the permitted purpose. Zero incidents in a small sample does not prove safety.
- Fairness: differences in failures and outcomes across relevant groups that may appropriately be measured, with uncertainty. Do not infer sensitive traits merely to populate a dashboard.
- Diversity: distributions of displayed and selected options and access to alternatives. Diversity does not automatically equal quality.

Safety and rights are independent constraints. A combined success score must not allow sales gains to offset their violation.

A falsifiable evaluation

H1: adaptation improves the customer outcome by a prespecified practically meaningful amount. H2: it improves contribution per eligible customer. H3: AI outperforms simple rules sufficiently to justify cost and risk. H4: agency and trust do not decline beyond a prespecified acceptable margin. H5: the effect persists beyond novelty. Some hypotheses can pass while others fail; do not report only favorable results.

Design A as a good standard offer, B as rules-based modules and C as AI-supported configuration/generation. With limited sample, begin with two groups and one question. Hold pricing policy and channel stable to isolate configuration effects. Randomize before exposure at user, account or branch level according to spillover risk. Analyze by original assignment, not just among people who adopt personalization.

Calculate sample size from baseline, variance, minimum meaningful effect and power. Thirty days is not automatically enough evidence. Check A/A where suitable, sample-ratio balance, lost events and identifier consistency. Prespecify the primary measure, guardrails and stopping criteria; address multiple comparisons. Do not stop at the first favorable p-value without an appropriate sequential design. Report absolute and relative effects, uncertainty intervals and costs. Nonsignificance does not establish equivalence or absence of harm.

Separate prespecified subgroup analysis from exploratory findings requiring confirmation. Shared inventory, staff or appointments can cause interference; use appropriate clustering or another suitable design. Track model changes and do not assume an old version's results apply to a new one. Microsoft research supports controlled experimentation and trustworthy analysis [3,4]; the application details here are a proposal to tailor to each setting.

A 30/60/90-day rollout

Days 1–30: choose one low-risk use case; research needs; establish a baseline, data map, module library and cost model; build simple rules; test in shadow mode and against failure scenarios; assign ownership and stop authority. Gate: real need, initial feasibility, reliable measurement and approved modules.



Days 31–60: run a limited experiment with clear choice, review and actual cost/exception tracking. Stop harm promptly according to the plan. Gate: sufficient evidence, not a calendar date. An insufficient sample means continued evidence collection rather than success.

Days 61–90: analyze effects, capacity pressure, fairness and trust. Scale gradually, simplify or stop. If AI and rules perform similarly, prefer the simpler option after considering cost and risk. Day 90 is a review point, not a promise of general validity or medical authorization.

Team worksheet

Complete one worksheet per use case, and revisit it before and after the experiment:

1. What is the use case, who is the customer, and what is the relevant context?
2. What outcome is intended, and what evidence establishes its importance to the customer?
3. Which data are strictly necessary? What purpose, permission and retention period applies to each?
4. What good standard alternative will be the comparator?
5. Which modules may be configured, and what must not change?
6. What does AI do, what does a person review, and who can reject the recommendation?
7. What are the cost, capacity and delivery-time limits?
8. How can the customer choose, edit, decline and leave?
9. What are the registered hypothesis, comparator and primary metric?
10. What are the guardrails and stopping thresholds, and how will breaches be detected?
11. What are the randomization unit, required sample and analysis plan?
12. Who owns delivery, monitoring and incident response?
13. When do preferences and inferences expire?
14. Is the decision to scale, modify or stop, and what evidence and uncertainty support it?

Failure modes

Cosmetic personalization: name the actual adaptation level. False precision: ask and refresh confidence.

Option explosion: reduce modules. Margin collapse: measure contribution per eligible person.

Undeliverable promises: validate capacity before offering. Lock-in and bias: review coverage, safe diversity and reset rights. Purpose violations: separate permissions and access. Manipulation: prohibit exploitation of fear and vulnerability. Nominal oversight: supply real time, authority and workload limits. Biased experiments: prespecify and analyze appropriately. Unjustified medical extrapolation: maintain clinical and regulatory separation. No desire for adaptation: respect the standard product as a legitimate outcome.

Primary references and their limits

[1] WHO, Ethics and governance of artificial intelligence for health, 28 June 2021. Supports health-AI governance, rights and safety, not the proposed marketing framework:

<https://www.who.int/publications/i/item/9789240029200>



[2] FDA, Clinical Decision Support Software, final guidance, January 2026. A specific US regulatory reference: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software>

[3] Kohavi et al., Online Experimentation at Microsoft, 2009. Controlled experimentation: <https://www.microsoft.com/en-us/research/publication/online-experimentation-at-microsoft/>

[4] Fabijan, Three Key Checklists and Remedies for Trustworthy Analysis of Online Controlled Experiments at Scale, 2019. Experimental quality and analysis: <https://www.microsoft.com/en-us/research/publication/three-key-checklists-and-remedies-for-trustworthy-analysis-of-online-controlled-experiments-at-scale/>

[5] NIST AI RMF 1.0, 2023, Core. Governance, mapping, measurement and risk management across the lifecycle, not sequential marketing stages or framework validation: <https://airc.nist.gov/airmf-resources/airmf/5-sec-core/>

The proposal is consistent with continuous governance in NIST [5], but is neither an official derivative nor endorsed by NIST. Its intended value is to connect understanding, configuration, delivery and evaluation in a reviewable practice, rather than claim to have invented personalization or abolished segmentation. Evidence from each application must determine its usefulness.