softworks labs

RESEARCH / 01

Egypt's digital choices

What the data says about access, trust and the cost of change



Ahmed Mostafa

People & technology

SWL-2026-01

4 November 2026 · v1.0

Abstract

Why do some digital services become part of daily life while others remain difficult to use? This working paper studies Egypt through access, trust, age, hospital workflows, business costs, security, reliability, and environmental limits. We analyze a frozen World Bank internet-use series, review source-scoped Egyptian and international research, and execute a local SQLite failure-and-retry benchmark. Egypt's internet-use value was 74.65% in the latest common comparison year, 2024. That measure does not explain motives or task skills. The reliability test applied 500 synthetic jobs: retries without an operation key created 50 duplicate effects, while transactional keyed retries created none under the same bounded trace. A separate 50,000-draw financial scenario explores chosen assumptions, not measured Egyptian returns. The evidence supports testing specific barriers and useful outcomes instead of treating Egyptians as one group that accepts or rejects technology. This paper publishes code, aggregate data, definitions, limits, and an editable decision model. It does not estimate national fear, rank every city, or predict complete adoption.

Inside the paper

1. The wrong question can give us the wrong product
2. How we built the evidence
3. Egypt is connected, but connection is only the beginning
4. The small words around a number matter
5. Age is a design question, not a judgement about people
6. What people may be protecting when they hesitate
7. A city map should not pretend to know more than the data
8. A hospital needs a working path, not just a new app
9. Scams use confusion as well as weak software
10. What public conversations add, and what they cannot add
11. Small firms face a different decision
12. What large-company evidence can teach a smaller team
13. Electricity changes the design, but does not explain every decision
14. Count the whole cost of a completed task
15. Why one precise saving is not enough
16. A small reliability experiment you can inspect
17. How to test a useful service in the real world
18. Useful technology should leave room for the physical world
19. A practical path for Egypt, one service at a time
20. Questions this paper does not answer
21. Read the files, change an assumption, rerun the result
22. Build something people can use with confidence

01

The wrong question can give us the wrong product

Do Egyptians dislike technology? It is an easy question to ask. It is also too wide to guide a useful product. A person may enjoy a payment app, avoid an online medical visit, and ask a family member to fill in a government form. Those choices can exist in the same life. Calling that person afraid of technology hides the work we need to do.

I want Softworks Labs to start with a better question: what makes a digital service worth trusting, learning, and using again? This paper looks at Egypt through that question. It joins public statistics, local studies, and small engineering experiments. It also leaves some questions open. A clear gap in the evidence is more useful than a confident story built on a guess.

Our main finding is simple. Internet use is not the same as useful digital adoption. A working connection helps, but people also need a clear price, an understandable process, support, and a safe way back when something goes wrong. For a business, the system must fit real work. For a patient, it must respect privacy and clinical responsibility. For an older person, it should allow learning without shame.

This is an applied working paper from Softworks Labs. We did not survey the Egyptian population or test patients. The new contribution is a checked public-data pipeline, a clear comparison of different kinds of evidence, a reproducible reliability experiment, and a practical way to test an automation decision.

02

How we built the evidence

The review includes sources available by 4 November 2026. We searched official statistics, university research, published studies, company case reports, and a small set of public discussions. This is a targeted review. It is not a systematic review of every paper, and it is not a meta-analysis. We did not combine unrelated studies into one national score.

Every important number needs four labels: what was measured, who was included, when it was observed, and where it came from. Publication year and observation year often differ. A report published in 2026 may describe 2025 work. A current website can still display an older survey. Our files preserve these differences instead of replacing every date with "today".

We used the World Bank API for national internet-use data. The underlying indicator comes from the International Telecommunication Union. The downloaded snapshot reports an update on 8 October 2026. We kept the raw response, calculated file hashes, and produced tidy tables with Python. A hash is a short digital fingerprint that helps identify an exact file. Our public package contains source notes, calculations, checks, and clearly marked synthetic results. D01

We gave different evidence different jobs. A national indicator describes broad access. A campus study helps us ask better hospital questions. A company success story shows a possible use, but does not prove a typical return. A public post can reveal a confusing step, but cannot tell us how common the problem is. This separation is central to the paper, not a small warning at the end.

03

Egypt is connected, but connection is only the beginning

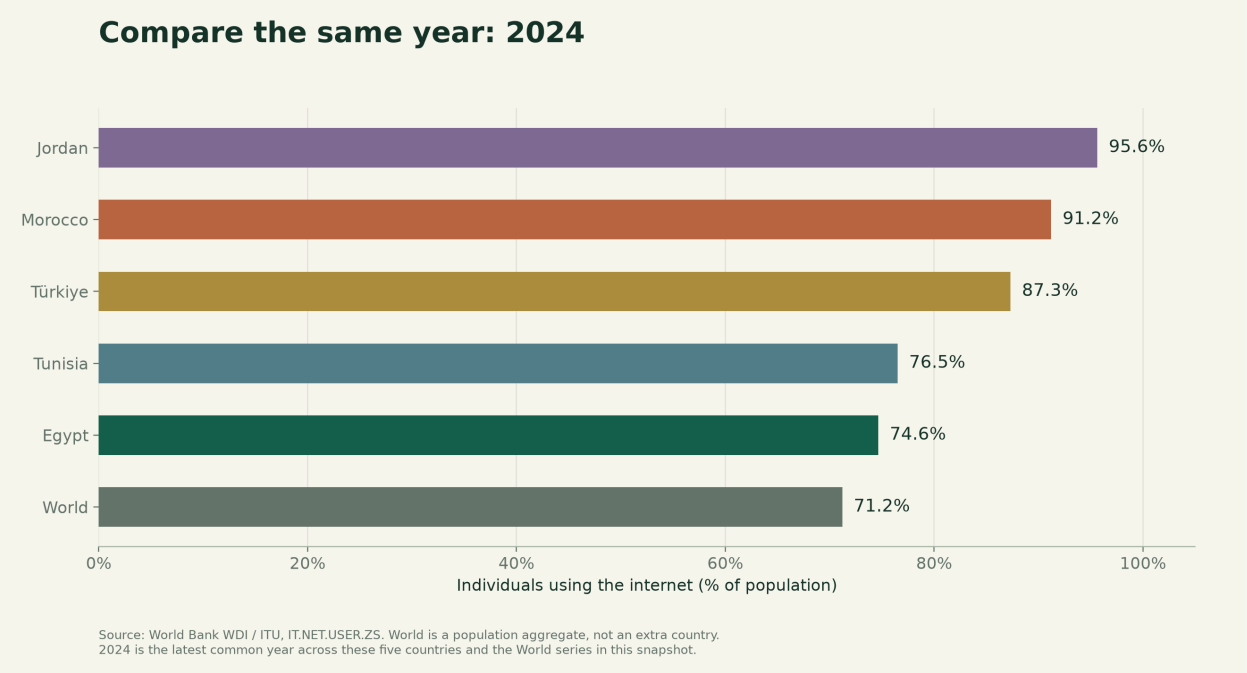
For a fair comparison, we selected the latest year available for Egypt and all four regional peers in this download: 2024. Internet use was 74.65% in Egypt, 76.50% in Tunisia, 87.31% in Türkiye, 91.20% in Morocco, and 95.62% in Jordan. The world value was 71.20%. These are shares of individuals who used the internet from any location in the last three months, not household subscriptions or shares using AI. D01

Egypt was therefore about 3.45 percentage points above the world value, but below these four peers. A percentage point is the direct difference between two percentages. Morocco's lead over Egypt was about 16.55 points. That is different from saying Morocco was 16.55% higher. The chart keeps the same year and measure across countries so that the comparison stays meaningful.

The long series is useful because a single year hides movement. Egypt rose from about 37.82% in 2015 to 74.65% in 2024, a gain of 36.83 percentage points. Its gap against the world moved from −2.08 points in 2015 to +11.81 in 2020 and +3.45 in 2024. Choosing only one of those years changes the story. These differences do not identify the effect of a policy or the pandemic. D01

Sources can revise earlier observations, and missing values do not mean zero use. We found an older mirror of the series with a different historical world value. We chose the original API snapshot and recorded that choice. A later reader may download a revised series and get a different result without either person making an arithmetic mistake.

These numbers cannot tell us how often someone uses the internet, whether their connection is reliable, or whether they can finish a form alone. A person who uses messaging once may be counted as an internet user. That does not make a complex business tool easy for them. The practical lesson is to test the actual task, not use national internet access as a substitute for usability research.



Internet use, common year 2024. Source: World Bank / ITU, frozen API snapshot. National individual-use measure; not service trust.

04

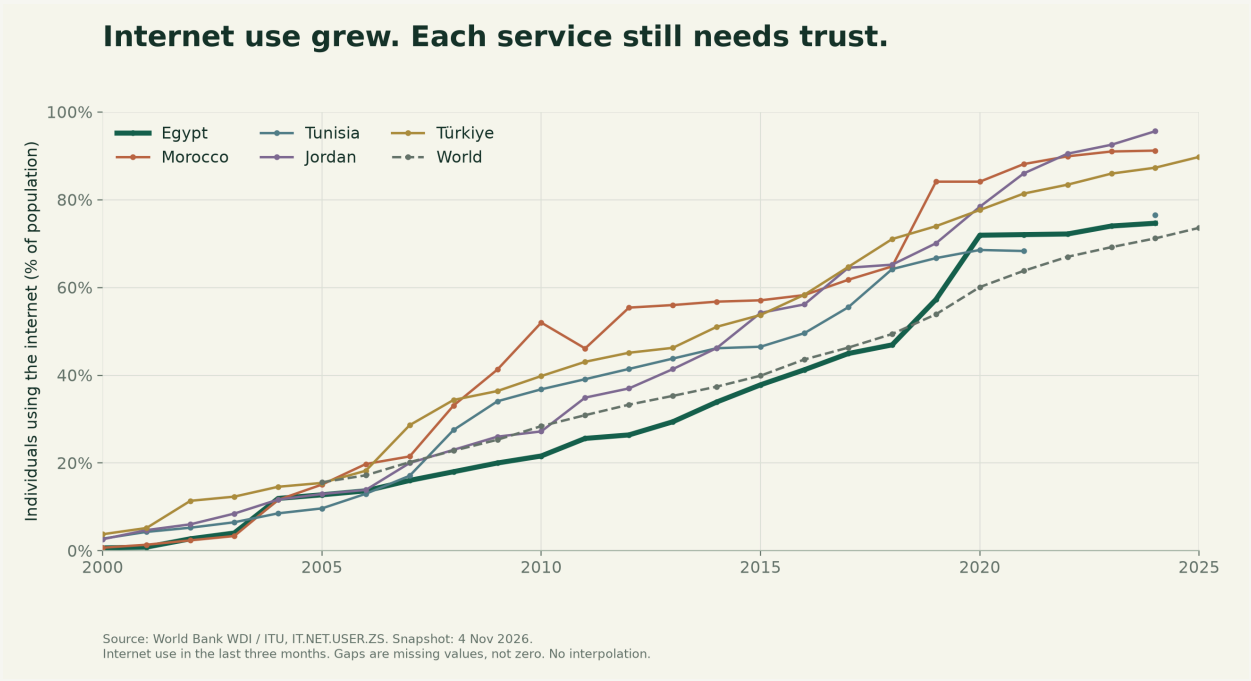
The small words around a number matter

Imagine two statements: “half the doctors use a service” and “half the people use a service”. They sound close. They describe different populations. The denominator is the group under the fraction. If that group changes, the meaning changes with it. Many weak technology arguments come from losing this small detail while moving a number from a paper to a slide.

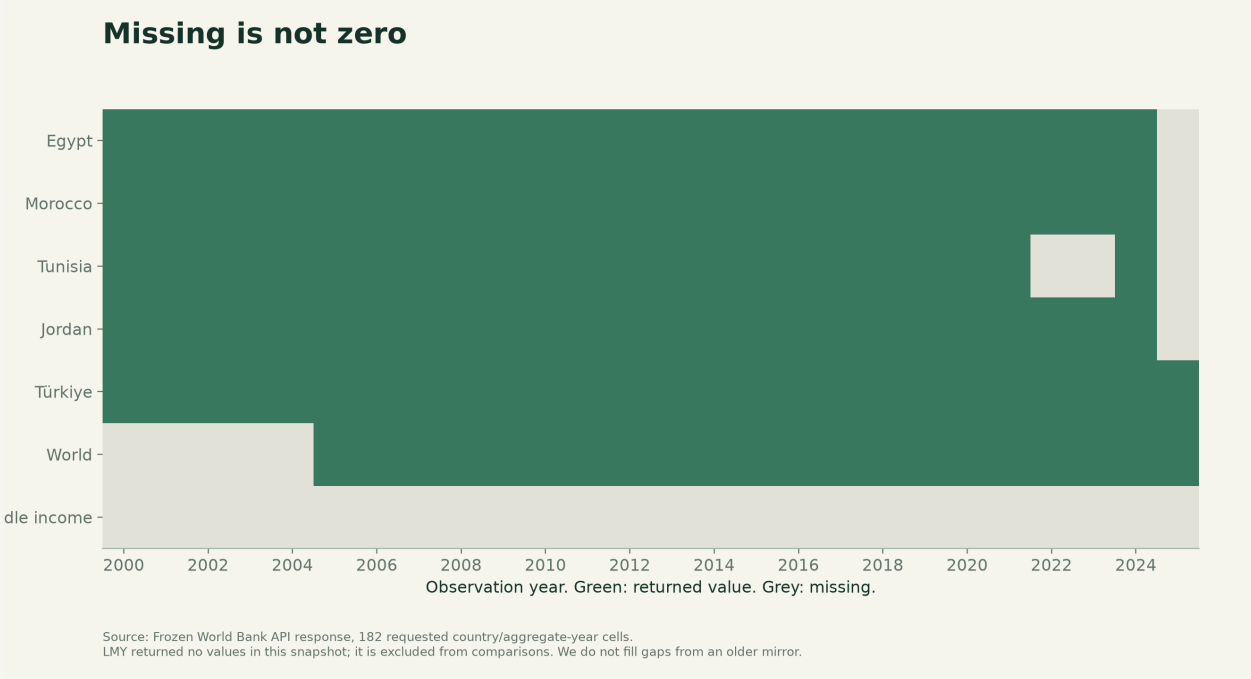
We use three rules. First, keep people, households, businesses, and recorded events separate. Second, keep a willingness answer separate from a completed action. Third, keep a percentage of users separate from a percentage of everyone. A feature can be popular among a small group of active users while still reaching very few people overall.

The same care applies to missing data. Our requested World Bank low- and middle-income aggregate, code LMY, returned no values in this snapshot. We did not fill that gap from an older mirror and present it as the same release. Tunisia also had gaps in the requested years. The public data keeps those cells empty. A line chart should not turn an unknown year into a measured point.

There is a useful test for any chart in a business meeting: can someone explain its unit, period, population, and source in one sentence? If not, the chart is not ready to support a decision. More decimal places will not fix that problem. We round values for reading, preserve the original precision in the data, and use direct differences only where the definitions match.



Observed internet-use series. Missing values remain missing; no interpolation. Same frozen World Bank / ITU snapshot.



Data availability in the requested API snapshot. Coverage is not adoption: an empty cell means no value returned.

05

Age is a design question, not a judgement about people

The ITU Egypt dashboard reports 2022 internet use of 93% for ages 15–24 and 65% for ages 25–74. It also reports household access of 84% in urban areas and 64% in rural areas. These are different measures with different denominators. Its oldest-age observation is from 2018, so we do not place it beside 2022 groups as a current age comparison. C01

A broad group such as 25–74 cannot explain the experience of a person aged 68. Age may overlap with income, schooling, work, health, device access, and help from other people. These data do not separate those pathways. We should not turn an age gap into a claim that older people are unable or unwilling to learn.

For product design, I would test practical barriers one at a time. Can the person read the screen without zooming every line? Does the button name match the action? Is there enough time to read a message before it disappears? Can a mistake be corrected? Can the person speak to a real helper without losing their place? These are testable questions. "Old people do not like apps" is not.

Family support can be valuable, but it can also hide dependence. A completed online form may mean that a younger relative did all the work. In a pilot, ask who completed each step and whether the user wanted that help. Offer a clear permission model for trusted support. Do not solve access by encouraging families to share passwords or quietly remove the older person's control.

06

What people may be protecting when they hesitate

Fear is one possible reason for refusal. It is not the only one. A person may protect money, privacy, time, dignity, or the ability to ask a question. If a service has failed before, caution may be a sensible response. Our evidence does not provide a national percentage for "fear of technology", and this paper does not create one.

A qualitative Egyptian study included 41 internet users and eight physicians. It found mixed feelings around online health information and the importance of communication with physicians. These are themes from a small selected group, not percentages of Egyptians. They suggest that useful digital activity and uncertainty can exist together. L05

The following is a design hypothesis, not a measured psychological model of Egypt. A person weighs an expected benefit against the effort to learn, the possible loss, and the difficulty of getting help. A small chance of a serious loss can matter more than a small saving in time. That is especially relevant when the person cannot easily replace lost money or repeat a long trip.

A useful interface therefore explains the next step before asking for commitment. Show the total price. Say what happens after submission. Give a receipt that can be found later. Explain who can see the information and how to correct it. Make cancellation and support visible. Test whether these changes improve completion

and confidence. Do not assume that a trust badge, a friendly illustration, or an AI label has solved the underlying problem.

07

A city map should not pretend to know more than the data

Egypt is not one user journey. A service may work differently across districts, workplaces, and households inside the same city. We found a 2023 survey covering Cairo, Alexandria, and Port Said. It interviewed 600 people aged 16–35 by telephone, 200 in each city, about household conditions. That is a narrow source for discussing city context, not a current national ranking of digital readiness. C02

During review, we found an inconsistency on page 53 between the narrative and chart for Alexandria's "seldom" response. We excluded the disputed cell. More importantly, we did not make a colourful map of every governorate from a three-city sample. An attractive map can make weak coverage look complete.

Our city panel therefore shows research coverage and practical questions. For Cairo, a pilot can test whether a service reduces repeated travel between offices. In Alexandria, it can test the handoff between online requests and physical service desks. In Port Said, it can test document status and support ownership in a trading workflow. These are proposed studies, not claims that those are each city's main problem.

For rural communities and Upper Egypt, the first task is to recruit people rather than extend assumptions from a capital-city sample. Include users with shared phones, limited data, low literacy, and unstable access. Record the place and conditions of each test. A local comparison becomes useful when the same task and method are repeated across places. Until then, the honest visual is a coverage diagram with clear gaps.

08

A hospital needs a working path, not just a new app

At Tanta University's medical campus, a convenience sample included 500 physicians and 500 employees in late 2023 and early 2024. Among physicians, 245 had used telemedicine. Of the 255 non-users, 156 wanted to try. Among the 245 users, 170 reported text messaging and 16 a personalised platform; multiple answers were allowed. This is one campus, not all Egyptian hospitals. L01

The result makes a simple "doctors reject technology" story difficult to defend. Some people were already using familiar tools. Others expressed interest without using them. That difference gives us a product question: what must change between interest and safe, repeated use? The study alone cannot tell us which intervention would work.

An AUC master's thesis used responses from 177 healthcare providers and interviews with seven experts. Its repository abstract identifies financial concerns, technical support, training, and data security as barriers. We verified the abstract, not the full restricted download, and do not assign a percentage to those themes. L02

A newer Mansoura study included 532 adult internet-using primary-care attendees. It classified 32.7% as having limited electronic health literacy under its chosen self-report scale and cutoff. It excluded non-users and people unable to read or write. Its mainly young sample cannot represent older Egyptians. L03

For a hospital pilot, begin with a bounded administrative task, such as explaining appointment requirements. Keep a named human responsible for exceptions. Measure wrong instructions, repeated calls, staff review time, and whether patients can reach help. Clinical decisions require a different validation process. Our small software experiment below is not clinical evidence and does not establish that an AI medical system is safe.

09

Scams use confusion as well as weak software

An official banking advisory warns against people requesting confidential banking information through calls or social media. It directs customers to official channels. This is prevention guidance, not a count of victims. It supports a clear design principle: a service should help users recognise the real channel before a stranger creates urgency. L09

On 23 October 2026, Egypt's Ministry of Interior reported arrests in a case involving allegedly fake discounted AI subscriptions. According to the police account, social pages collected advance payments and then cut contact. We did not verify a court outcome. The case illustrates a pattern; it does not measure how common that pattern is. L10

INTERPOL reported 12,281 ransomware detections attributed to Egypt in 2024 in Trend Micro data. A detection is not a unique victim, a successful breach, or a hacked company. Coverage depends on the vendor's systems. We do not divide that number by Egypt's population to create an individual risk rate. L11

Our practical response is to reduce ambiguity. A payment screen should make the payee, amount, and result clear. Support should never need a user to reveal a one-time code. Account recovery needs careful checks because it can become a second entrance to an account. Staff need a safe way to pause a suspicious request without pressure to finish it quickly.

For users, the simple rule is to stop, leave the message, and contact the institution through a known official route. For builders, the harder job is to make that route easy to find and make help responsive. Blaming a person after a confusing failure teaches the next person to hide a problem. Good recovery is part of adoption.

10

What public conversations add, and what they cannot add

We reviewed five selected public Reddit discussions. They included concern about unfamiliar card payments, interest in automatic expense tracking, questions about paid-app value, a wish to connect services, and confusion over which organisation should solve a payment problem. These are unverified personal accounts. We retained no handles, private details, or claimed national sentiment score. F01 F02 F03 F04 F05

The selection is small, finance-focused, and shaped by search results. People who post online already differ from many people who do not. A thread full of complaints can reflect where people seek help, not the average service experience. We found no usable no-login X sample, so there is no invented analysis of Egyptian tweets.

The useful output is a list of research questions. Does a customer know who owns a failed transfer? Can they tell whether an action is pending or complete? Do existing apps connect well enough for everyday work? Is a subscription's extra value clear? A product team can test those questions directly with recruited users.

This is also why we did not run a sentiment model and publish a negative-percentage headline. A model can count labels on a selected set of posts very precisely while answering the wrong population question. Better analysis starts with the sampling problem. More computing power cannot repair a dataset that never included the people we want to understand.

11

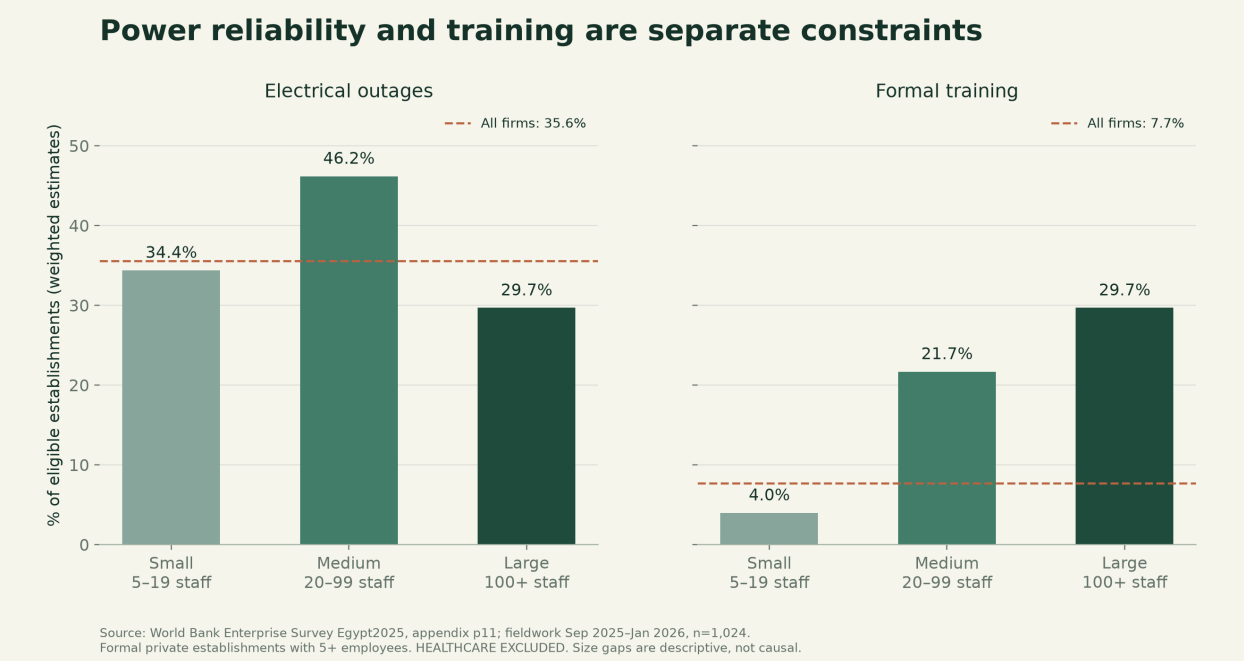
Small firms face a different decision

The World Bank's Egypt 2025 Enterprise Survey interviewed 1,024 establishments between September 2025 and January 2026. It covers eligible formal private businesses with at least five employees. It excludes several sectors, including healthcare and education, and does not represent every microbusiness or informal seller. Its published weighted estimates show that 35.6% experienced electrical outages and 7.7% offered formal training. Weighting adjusts how much each sampled firm contributes to the estimate. D02

Those two figures describe different conditions. They do not show that lack of training caused outages or that either caused resistance to AI. The size breakdown also resists a simple story: reported outage exposure was 46.2% for medium firms, compared with 34.4% for small and 29.7% for large firms. Bigger was not always better in this measure. D02

A separate ERF working paper using 2022 firm data reports website use of 32% among small Egyptian firms and 84% among large firms. We keep its size definitions separate from the World Bank survey. A website is also a narrow measure of digitisation; a seller may use messaging or payments without owning one. L06

A small company has to pay for setup, learning, maintenance, and mistakes. Those costs arrive before a promised saving becomes real. If work volume is low, a manual step can remain cheaper. If a process changes every week, automation can move the work into debugging. The useful question is not whether the owner is modern. It is whether the new process produces better completed work at a sustainable total cost.



Selected weighted firm estimates, Egypt 2025 survey; fieldwork September 2025–January 2026. Small: 5–19 employees; medium: 20–99; large: 100+. Not hospitals or informal microbusinesses.

12

What large-company evidence can teach a smaller team

The final published version of "Generative AI at Work" studied 5,172 customer-support agents at one company. Access to an AI assistant increased resolved issues per hour by about 15% on average, with different effects across workers. It was a workplace rollout study, not an estimate of Egyptian company returns. L12

A separate randomised experiment with 453 professionals found that assisted writing tasks took 40% less time and received 18% higher quality scores. These were bounded writing tasks. Neither figure means a whole company will gain the same amount in revenue or profit. L13

There is also a useful counterexample. METR's study of early-2025 tools involved 16 experienced developers and 246 tasks in familiar repositories. AI-allowed work took 19% longer in that setting. It does not describe all coding or the tools available in November 2026. Together, these studies show why we should measure the task and context instead of promising one universal AI gain. L18

In Egypt, a B.TECH and McKinsey case report claims a 4% gross-profit uplift from AI pricing. It separately reports a 14% conversion uplift from e-commerce changes. These are client and consultancy claims with limited public methods, not independently verified effects from our research. We do not add the two numbers together. L15

A small team can copy the discipline of a bounded test without copying a large company's budget. Select a repeatable task, define acceptable quality, log review and repair time, and compare with the existing process. A faster first draft is useful only if the final work is correct enough to use.

13

Electricity changes the design, but does not explain every decision

An outage can stop a local machine, a router, or a work session. A cloud service may remain online while the employee cannot reach it. This creates a gap between a provider's uptime and the user's ability to finish work. A useful service needs to think about both ends of the connection.

The enterprise survey establishes reported exposure among its eligible firms. It does not identify how much foreign investment was lost because of power reliability. We found no defensible causal estimate for that question in this review. Prices, demand, financing, policy, skills, and risk can all enter an investment decision. Assigning the whole decision to electricity would be a stronger claim than the evidence allows.

Our engineering response is a proposal: save progress early, keep a clear pending state, use a durable queue where appropriate, and make repeated requests safe. If a connection fails after a request reaches the server, the user may not know whether it worked. Pressing again should not create a second bill or booking.

This is where idempotency helps. It means that repeating the same intended action does not repeat its side effect. The server keeps a record of the action's unique key and returns the earlier result when the request is repeated. A key alone is not enough. The key and the action need correct transaction handling, suitable retention, and checks that the repeated request has the same meaning.

A fallback should also be tested as a real service. Paper or a phone call is not automatically reliable. Decide how offline work returns to the main record, who resolves conflicts, and how users know which state is final. Continuity is a workflow, not simply a backup generator or another button.

14

Count the whole cost of a completed task

Consider a hypothetical company processing 2,000 requests a month. Manual work takes eight minutes each. That is 16,000 minutes, or about 266.7 hours. At an assumed loaded labour cost of EGP 60 per hour, the time costs EGP 16,000. "Loaded" means the full cost attached to the work, not a claim about a typical

Egyptian wage.

Now suppose automation still needs 1.5 minutes of handling and two minutes of review per request. If 5% of requests need ten extra minutes of repair, the average total becomes four minutes. That releases about 133.3 hours, valued at EGP 8,000. Subtract an assumed EGP 2,500 monthly service cost and EGP 800 of additional disruption cost: four hours at EGP 200 per affected hour. The modelled net monthly value is EGP 4,700.

The equation is: net monthly value = jobs × (manual minutes − automated handling minutes − review minutes − error share × repair minutes) ÷ 60 × hourly cost − service cost − additional disruption cost. If setup costs EGP 12,000, simple payback is about 2.55 months at EGP 4,700 per month. These are the calculator's defaults: an illustration, not a forecast, market price, or measured client result. Count disruption only when it is additional; do not also deduct the same lost work from job volume.

Saved hours are not automatically cash in a bank account. A business may use them to reduce queues, improve service, or handle more work with the same team. Revenue only rises if there is demand and the company can deliver. Our calculator keeps these assumptions visible so the reader can change them rather than accept a single attractive ROI figure.

Quality should be a constraint, not an afterthought. A cheap workflow that sends the wrong invoice may be more expensive than the old one. For each task, count accepted completions, corrections, unresolved cases, and the time of everyone involved. Decide what level of error makes a system unacceptable before seeing the result.

15

Why one precise saving is not enough

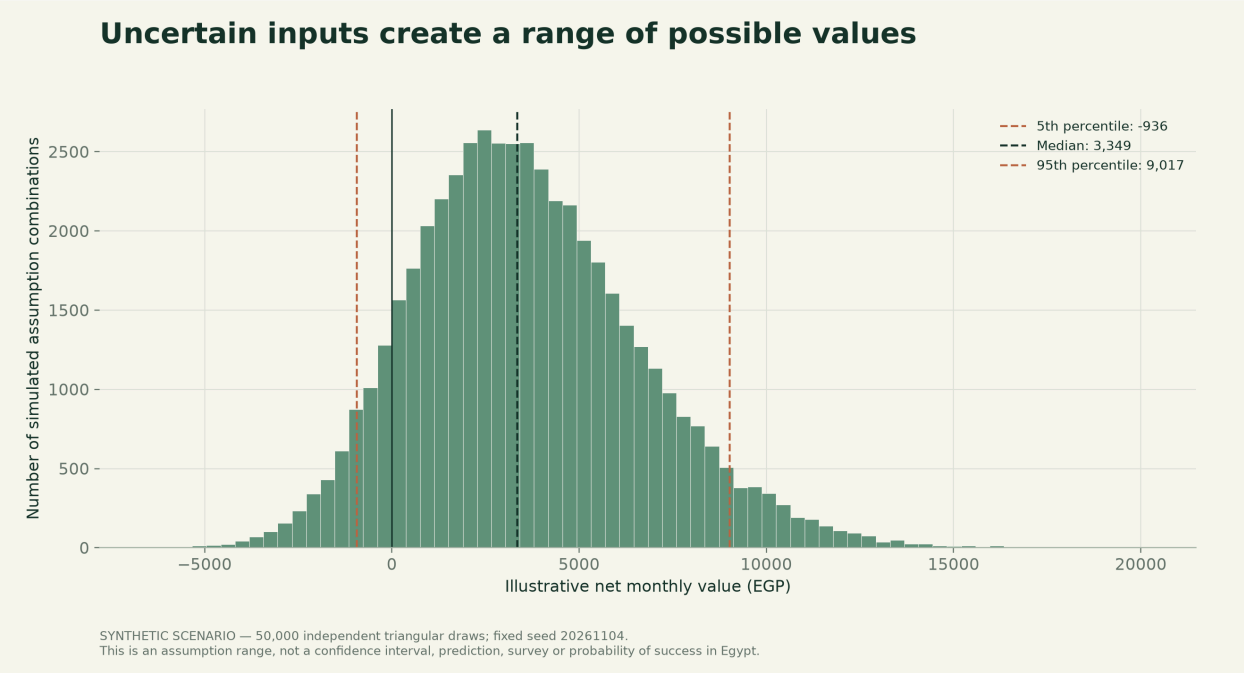
The example above uses fixed assumptions. Real work moves. Volume changes, review takes longer on difficult cases, and the failure rate may rise after a supplier changes a format. The accompanying Python analysis uses a seeded Monte Carlo scenario to show how uncertainty in inputs changes the output. Monte Carlo means repeating a calculation with many sampled combinations of assumptions.

We ran 50,000 draws with seed 20261104. Modelled monthly value had a median of about EGP 3,349; the 5th and 95th percentiles were about −EGP 936 and EGP 9,017. About 89.52% of these assumed scenarios were positive. Inputs use independent triangular distributions with stated minimum, most likely, and maximum values; the complete ranges are in the JSON.

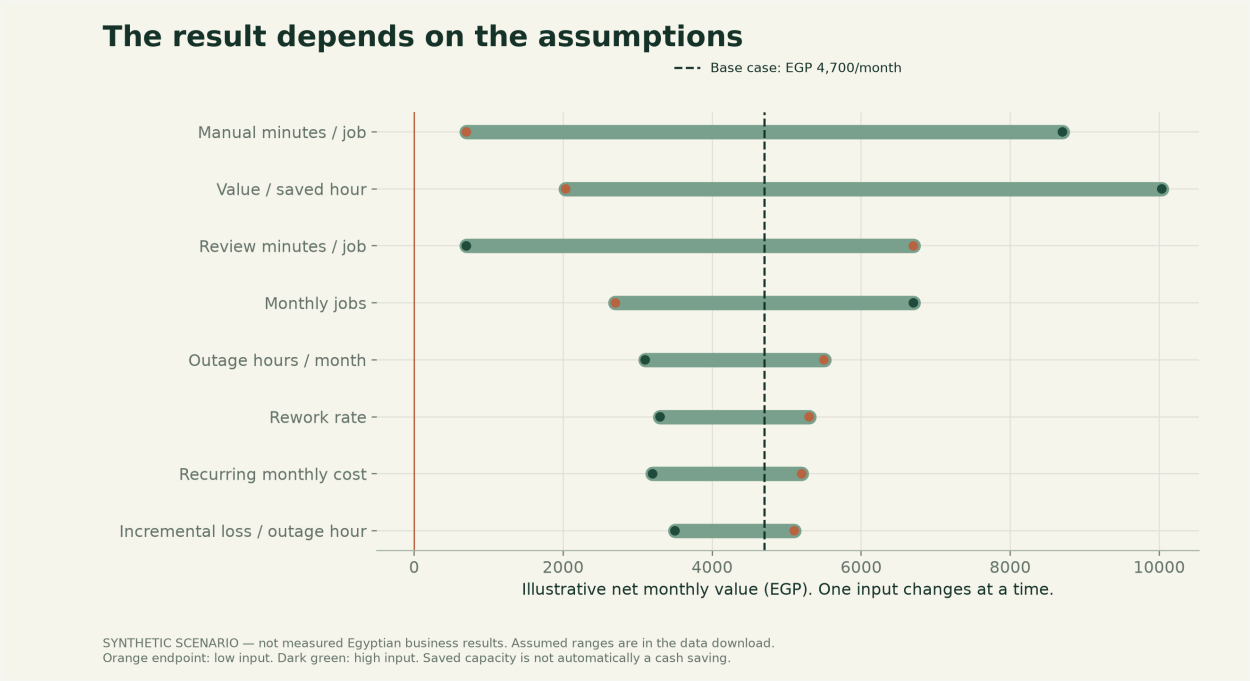
This is synthetic analysis. The input ranges are not estimates from Egyptian companies. The output distribution is not a confidence interval for the country, and its positive share is not the chance that your business will succeed. It answers a narrower question: given the chosen model and ranges, how often does this calculation produce a positive monthly value?

The seed makes the run repeatable. Someone using the same code, input settings, and supported environment should get the same simulated draws. Reproducibility does not make the assumptions true. It makes them visible and easier to challenge. That distinction is one reason we publish both the code and the settings.

Before a real purchase, replace the invented ranges with observations from a short pilot. Measure the full distribution of review time, not only its average. Keep difficult cases. Test whether failure and review time rise together; our simple independent-input example does not model all such relationships. Repeat the calculation under a lower work volume and higher repair cost. A decision that only looks good under perfect conditions needs more work.



Synthetic model: 50,000 draws, seed 20261104. Chosen independent assumptions; not observed Egyptian returns or a statistical confidence interval.



One input changes at a time while the others stay at default values. The ranges are chosen assumptions, not measured uncertainty.

16

A small reliability experiment you can inspect

We built and executed a local SQLite fault-injection benchmark around repeated requests. Fault injection means deliberately creating a controlled failure so that a recovery method can be tested. We compared one-shot delivery, retries without a key, and retries with a transactional operation key. Each strategy received 500 synthetic jobs in each of 20 measured trials after one warm-up. Execution order rotated. No live payment service, patient record, customer account, or production network was involved.

The fixed trace included 125 jobs whose first attempt failed before writing and 50 configured to lose their first confirmation after a successful write; 15 jobs were in both groups. One-shot delivery applied 375 unique jobs, confirmed only 340 to the caller, and left 125 unapplied. Retrying without a key applied all 500 but created 50 duplicate effects. Transactional keyed retries applied all 500 with zero duplicate effects. Both retry paths made 675 attempts. These are measured software outcomes under a deliberately chosen trace, not estimates of real outage frequency.

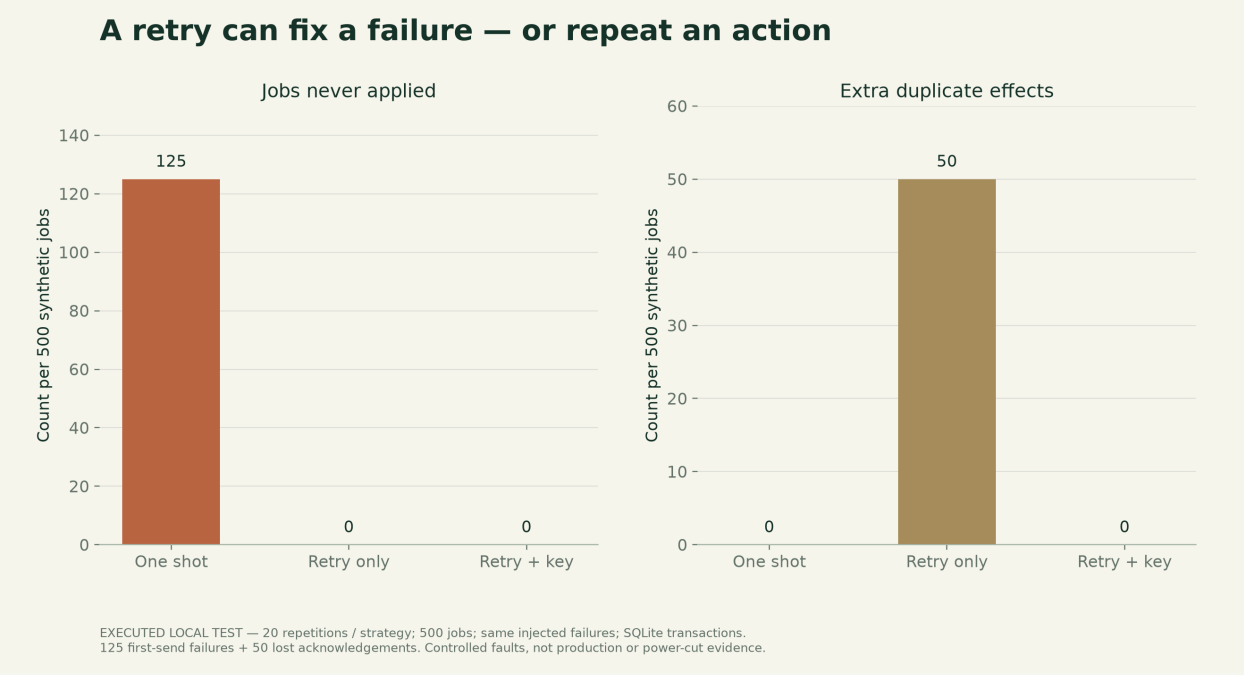
Median whole-trial elapsed times were about 410 ms, 602 ms, and 585 ms respectively. They describe local SQLite work, not network latency or a proven speed advantage. A separate outbox test closed and reopened database connections twice and finished with 500 unique actions, no duplicates, and no pending jobs. That was not an operating-system crash or a physical power-cut test. Persistent outages, concurrent writers, different keys for the same human intent, external side effects, and key expiry remain outside the test.

The browser example illustrates the same recovery idea in memory; it is not the SQLite benchmark. The Python code and recorded results are the reproducible experiment. Its three-attempt retry budget works for this bounded failure trace. It cannot guarantee delivery through an outage that continues beyond the budget. Production design needs durable queues, backoff, monitoring, and an explicit path for unresolved work.

This is the core transaction from the tested Python implementation. The same transaction records the operation key and its effect. The full file includes the setup, injected failures, retry loop, and assertions.

```
with self.con:
    added = self.con.execute(
        'INSERT OR IGNORE INTO idempotency(job_id) VALUES (?)',
        (job_id,),
    ).rowcount
    if added:
        self.con.execute(
            'INSERT INTO effects(job_id) VALUES (?)', (job_id,)
        )
```

The result matters because it connects a human concern to an engineering choice. "I pressed it, but I do not know what happened" is a trust problem with a technical part. A clear status screen and safe recovery can address that part. A cinematic interface alone cannot. Our goal is to make the failure understandable and the recovery inspectable, rather than hide them behind a successful animation.



Executed local SQLite benchmark, 500 synthetic jobs per trial; 20 measured trials per strategy. Identical injected failure trace. Delivery counts are deterministic; this is not a real outage-frequency estimate.

Strategy	Unique applied	Duplicate effects	Attempts
One shot	375	0	500
Retry without key	500	50	675
Retry with key	500	0	675

17

How to test a useful service in the real world

For a first field pilot, I would start with one low-risk workflow and a written stop rule. For example, a business could test document-status messages while people keep control of approvals. The team should agree on the task, users, acceptable quality, support path, and observation period before the test begins.

Where practical, randomly assign comparable tasks or users to the existing and proposed process. If random assignment is not possible, record why and describe the comparison as observational. Keep the time period, workload, and rules as similar as possible. A before-and-after improvement may reflect quieter days, more experienced staff, or a changed customer mix.

Choose one main outcome: perhaps total staff minutes per correctly completed request. Also track wrong outputs, repeated visits, unresolved cases, user effort, and help requests. Record the time spent setting up and maintaining the system. An improvement that quietly moves work from staff to customers is not necessarily a service improvement.

Recruit beyond the easiest users. Include people who need help reading, people with older devices, and people who choose a human route. Report how many people were invited, how many joined, and how many left. Explain missing records. Do not publish a strong success percentage after dropping every failed session.

The sample size should follow the smallest improvement worth detecting and the variation in real tasks. This paper cannot supply a universal number of participants. Run a small feasibility phase to learn about variation, then plan the comparison. Protect personal data, obtain appropriate consent, and seek the required oversight for any health study. Our proposed administrative pilot does not replace those requirements.

18

Useful technology should leave room for the physical world

The IEA estimates that data centres used about 485 TWh of electricity globally in 2025 and projects about 950 TWh in 2030 in its central outlook. These numbers cover all data centres, not AI alone. The future value is a projection, not a measured result. They cannot tell us the energy cost of one Egyptian app. L16

The Global E-waste Monitor reports 62 million tonnes of discarded electrical and electronic equipment in 2022, with 22.3% documented as formally collected and recycled. It includes many types of equipment. The remaining share is not automatically a measure of landfill waste. L17

Our design proposal is modest: use a large model only where it adds enough value. A rule, a database search, a small model, or a clear form may solve the task. Cache safe repeated results when the data and permissions allow it. Avoid sending the same work through several services without a reason. Support older devices and slow connections when possible.

Measure before making a green claim. A process may reduce travel but increase computing and hardware use. Define the boundary: which trips, devices, energy, and period are included? If a provider does not disclose enough information, say that the footprint is unknown. Do not turn fewer clicks into an invented carbon-saving number.

Nature is not a decoration around this argument. It is a reminder that systems use land, materials, electricity, and human attention. Progress can include repair, longer device life, less repeated work, and a service that lets someone finish quickly and return to life outside the screen. More technology is not, by itself, the outcome we should measure.

19

A practical path for Egypt, one service at a time

There is no evidence-based date when Egypt will become "fully evolved". People are not a product version. A country can improve access and services while keeping cash, face-to-face care, paper backups, and personal choice where they remain useful. The better target is that people can complete important tasks safely, with less effort and fair access.

For a city team, begin with one shared service problem and a public measurement plan. Use the same task definition across participating districts. Publish where the sample came from and where it did not. Compare completion, time, support, and failure recovery. Avoid a league table that punishes a district for serving harder cases.

For a hospital, start with clear responsibility and safe administrative boundaries. For a small business, start with repeated work and total cost. For a bank or payment service, start with clear status, recognisable support, and recoverable mistakes. For an education team, test whether the person can perform a useful action independently after training, not whether they watched a tutorial.

For older users, give time, readable language, and control over assistance. For workers, explain what changes in their job and listen to the work that the system cannot see. For managers, reward early reporting of failures. If people fear punishment for finding a problem, the dashboard can look healthy while the service becomes worse.

The first stage is to observe. The second is a bounded pilot. The third is a measured expansion with a rollback plan. The fourth is continued review after the exciting launch period. These are proposed steps, not a claim that one method fits every institution. Each stage should leave evidence that another team can inspect.

20

Questions this paper does not answer

We cannot estimate the national share of Egyptians who fear technology. We cannot rank all cities by digital trust. We cannot infer older people's attitudes from a young hospital sample. We cannot calculate foreign investment lost to electricity outages from the share of firms reporting them. We cannot promise that an AI system will save money in a company we have not measured.

The national data are broad and may be revised. The hospital studies have local samples and different methods. The enterprise survey excludes important parts of the economy. Company case reports may select successful projects. Public discussions have strong selection bias. The reliability benchmark and financial scenario have deliberately limited conditions. These limits shape the conclusions rather than cancel the value of the work.

The most useful unanswered question is how these barriers combine in one real journey. Does clearer recovery matter more than faster loading for a particular service? Does staff training help only when support is available? Do users keep control when a relative assists them? A well-designed local pilot can answer some of these questions. A large general article cannot answer them by confidence alone.

We also do not claim this is the first work to connect access, trust, and design. Our added value is practical: putting definitions beside numbers, exposing gaps, testing one failure pattern, and publishing a reusable decision model. The next version should improve through new data and documented corrections, not stronger adjectives.

21

Read the files, change an assumption, rerun the result

The download package contains machine-readable JSON and CSV, source and licence notes, Python acquisition and analysis code, and the recorded outputs. The public World Bank series is attributed under its data terms. Restricted source PDFs and personal social-media data are not repackaged as an open dataset. Short cited discussions remain linked to their originals.

Use a clean Python environment and the package instructions. Run the validation step before using the results. The acquisition script preserves supplied files by default and fetches the source PDF excluded from the ZIP so it can verify the survey table. Only its explicit `--refresh` option replaces the data snapshot. A new snapshot can contain revisions, so give a changed analysis a new date and keep both manifests.

The online charts have readable values and explanations. The calculator is a scenario, and the retry demonstration is a local model. Neither asks for personal records. The English and Egyptian Arabic editions use the same numerical findings. The PDFs preserve references and paper structure for readers who want to study or share a fixed version.

Readers should be able to challenge a number without having to trust the tool that wrote the sentence around it. Each downloadable result has a definition and a path back to its source or calculation. The fixed edition helps another team identify exactly what it is checking.

22

Build something people can use with confidence

My conclusion is not that Egyptians need to be pushed into technology faster. It is that useful services need to earn their place in daily life. Access matters. So do cost, support, control, reliable work, and the ability to recover from a mistake. A person's hesitation can point to a design problem worth solving.

For Softworks, that means starting with the work itself. What is repeated? What fails? Who pays for the failure? What should remain a human decision? Then we can choose a small experiment, measure a real outcome, and build the next step from evidence. Sometimes that will include AI. Sometimes a simpler tool will be the better answer.

This paper is an invitation to make those decisions more visible. Share the source, the definition, the failure, and the correction. Keep the person using the service inside the story. Keep the physical world inside the cost. Make progress something we can explain and inspect.

A small glossary

adoption
Starting to use a service and continuing to use it in real work.

denominator
The whole group under a percentage or fraction.

percentage points
The direct difference between two percentages: 60% minus 40% is 20 points.

sample
The people or organisations actually included in a study.

selection bias
A distortion when the people included differ in relevant ways from those left out.

causal
About whether changing one thing actually changes another.

confidence interval
A range from a statistical method describing uncertainty in an estimate; not the same as our scenario range.

Monte Carlo
Repeated calculations using sampled assumptions to explore possible model outputs.

idempotency
A property where repeating the same intended operation does not repeat its effect.

transaction
A group of database changes that succeeds together or is rolled back together.

fault injection
Creating a controlled failure on purpose to test recovery.

benchmark
A repeatable test used to compare methods under stated conditions.

synthetic
Created for an example or experiment; not observations from real customers.

telemedicine
Healthcare communication or services provided across distance using technology.

median
The middle value after the results are placed in order.

percentile
A position in ordered results; the 95th percentile has about 95% of values at or below it.

outbox
A stored list of intended actions waiting to be sent or confirmed.

hash
A digital fingerprint used to check whether a file has changed.

peer review
Independent specialists assess research methods and claims; our AI editorial checks are not this.

rollback
Returning to an earlier working state when a change causes problems.

References and source notes

[D01] World Bank / ITU — Individuals using the Internet (% of population), IT.NET.USER.ZS
API snapshot retrieved 4 November 2026; last updated 8 October 2026. National individual-use measure. Common comparison year: 2024. CC BY 4.0 data terms.
https://data.worldbank.org/indicator/IT.NET.USER.ZS

[D02] World Bank Enterprise Surveys — Egypt 2025 Country Profile
1,024 eligible formal private establishments, fieldwork September 2025–January 2026. Healthcare and informal microbusinesses are not covered. Weighted published estimates.
https://www.enterprisesurveys.org/content/dam/enterprisesurveys/documents/country/Egypt-2025.pdf

[C01] ITU — Egypt Digital Development Dashboard
Dashboard frozen through December 2023. Age and location observations quoted are from 2022; oldest-age data are from 2018. Mixed years and rounded values. Original dashboard not redistributed.
https://www.itu.int/en/ITU-D/Statistics/Documents/DDD/ddd_EGY.pdf

[C02] Austrian Federal Office for Immigration and Asylum — Egypt Socio-Economic Survey 2023
Published 2024. Telephone interviews, 4–22 December 2023; 600 respondents aged 16–35, 200 each in Cairo, Alexandria and Port Said. Household questions. Disputed narrative cell excluded; no national city ranking.
https://www.ecoi.net/en/file/local/2106870/2024-04-03_Dossier%2BEgypt%2BSocio-Economic%2BSurvey%2B2023.pdf

[L01] Knowledge, attitude, and practice regarding telemedicine among physicians and employees at Tanta University Hospitals, Egypt
Eman A. Younis; Amira K. El-Shenawy; Sanaa A. E. Abdo Published 2025-06. Tanta University medical campus; convenience sample of 500 physicians and 500 employees peer-reviewed cross-sectional survey Employees are not a patient sample. Convenience sampling and one campus prevent national inference. Willingness is not completed adoption. Reported age difficulties are perceptions, not a trial of older patients.
https://pmc.ncbi.nlm.nih.gov/articles/PMC12174010/

[L02] The User-Side Perspective on Implementing Electronic Health Records in Egypt
Nayrouz Ibrahim A. Romana Published 2024. 177 healthcare-provider questionnaire respondents and seven healthcare-management experts AUC master's thesis; mixed methods Thesis, not a journal article. Recruitment/generalizability and effect sizes were not verified from full text; download returned403. Do not turn listed factors into prevalence or causal effects.
https://fount.aucegypt.edu/etds/2315/

[L03] Assessment of electronic health literacy and its associated factors among internet-using primary health care attendees in Mansoura district, Egypt: a cross-sectional study
Shymaa Mamdouh Mohamed Abdu; Sahar Mohamed Youness; Abdel-Hady El-Gilany Published 2026-08-19. 532 adult internet-using PHC attendees in ten urban/rural facilities in Mansoura district peer-reviewed cross-sectional survey Excluded people who could not read/write and non-internet users. Predominantly young sample, not all patients or all Egyptians. eHEALS measures perceived skills; cutoff not universal. Associations cannot establish causes.
https://pubmed.ncbi.nlm.nih.gov/42618911/

[L05] The Attitudes of Egyptian Web-Based Health Information Seekers Toward Health Information Provided Through the Internet: Qualitative Study
Mayada Ghweeba and colleagues Published 2022-02-18. 49 participants:41 Egyptian internet users and8 physicians peer-reviewed qualitative study Purposive qualitative themes are not percentages of Egyptians. Participants already sought information online. Does not prove online information improves health.
https://formative.jmir.org/2022/2/e30108

[L06] Which Firms Are More Digitized? A Comparative Study between Egypt and Jordan
Chahir Zaki Published 2023-04. Firm Digitalization survey; paper describes around1,000 observations per country, across two waves ERF working paper1635; weighted firm survey analysis Do not import World Bank employee cutoffs into these size categories. Paper explicitly warns associations are noncausal. Pandemic-era data; exact subgroup sample counts/cutoffs need microdata documentation. Public PDF is copyrighted; cite/link, do not mirror publicly.
https://erf.org.eg/app/uploads/2023/04/1682432769_852_760052_1635.pdf

[L09] Together we fight fraud / معاً نكافح الاحتيال
Central Bank of Egypt statement republished by Egyptian Banking Institute, Masrafeyoun Published 2025-09-25. Bank customers; no statistical sample official banking-institute advisory Prevention guidance, not evidence of fraud prevalence or program effectiveness.
https://masrafeyoun.ebi.gov.eg/2025/09/25/%D9%85%D8%B9%D8%A7%D9%8B-%D9%86%D9%83%D8%A7%D9%81%D8%AD-%D8%A7%D9%84%D8%A7%D8%AD%D8%AA%D9%8A%D8%A7%D9%84/

[L10] Arrest of two people accused of defrauding citizens seeking AI-platform subscriptions
Egyptian Ministry of Interior Published 2026-10-23. One reported investigation; two people arrested official police incident announcement Describe as a reported case/allegation, not a conviction. It shows a scam pattern, not its national frequency or typical victim age.
https://moi.gov.eg/news/GetDetails?newsId=4fd5475b-cde3-46f3-9da5-ccc2e534997e

[L11] New INTERPOL report warns of sharp rise in cybercrime in Africa
INTERPOL Published 2025-06-23. Trend Micro-observed activity; coverage not a random sample of Egypt official report summary using law-enforcement and vendor data Detections are not unique victims, confirmed breaches, attacks that succeeded or population risk. Vendor coverage affects country comparisons. Never label12,281 hacked Egyptian companies.
https://www.interpol.int/en/News-and-Events/News/2025/New-INTERPOL-report-warns-of-sharp-rise-in-cybercrime-in-Africa

[L12] Generative AI at Work
Erik Brynjolfsson; Danielle Li; Lindsey Raymond Published 2025-02-04. 5,172 customer-support agents at one Fortune500 business-software firm peer-reviewed workplace quasi-experiment One workplace, historical GPT-3 assistant. Does not establish Egyptian ROI, layoffs, wage changes or universal gains. Earlier draft14%/5,179 differs; use published15%/5,172 consistently.
https://academic.oup.com/qje/article/140/2/889/7990658

[L13] Experimental evidence on the productivity effects of generative artificial intelligence
Shakked Noy; Whitney Zhang Published 2023-07-13. 453 college-educated professionals doing incentivized occupation-specific writing tasks peer-reviewed randomized experiment Short professional writing tasks, not all work or Egyptian companies. Time decrease is not the same as revenue increase. Early444-person working paper is not the final453-person version.
https://doi.org/10.1126/science.adh2586

[L15] How B.TECH transformed into a technology-powered retailer
McKinsey with client B.TECH Publication date not verified; accessed 4 November 2026. One Egyptian retailer and transformation program consultancy/client self-report Undated page; publication day not verified. Self-reported, selected success case, not independently audited. Conversion features and AI pricing are distinct interventions. EBITDA improvements cannot all be attributed to AI.
https://www.mckinsey.com/industries/consumer-packaged-goods/how-we-help-clients/how-b-tech-transformed-into-a-technology-powered-retailer

[L16] Key Questions on Energy and AI
International Energy Agency Published 2026-04-16. Global data centres institutional energy estimates and scenario projections All data centres, not AI-only.2030 is a projection, not observed fact or certainty. Global demand cannot be assigned to one Egyptian application.
https://www.iea.org/reports/key-questions-on-energy-and-ai/executive-summary

[L17] The Global E-waste Monitor 2024
UNITAR and ITU Published 2024. Global discarded electrical/electronic equipment international statistical assessment Includes many equipment categories beyond computers. Undocumented does not mean every remaining tonne went to landfill. Not Egypt-specific quantities.
https://www.itu.int/pub/D-GEN-E_WASTE.01-2024

[L18] Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity
METR; Joel Becker, Nate Rush, Beth Barnes, David Rein Published 2025-07-10. 16 experienced open-source developers;246 real tasks in familiar mature repositories randomized controlled trial; research report/preprint Small selected developer sample and early2025 tools. Does not show all AI hurts productivity or describe November2026 performance. See later METR follow-up separately; do not mix versions.
https://metr.org/Early_2025_AI_Experienced_OS_Devs_Study-paper.pdf

[F01] Debit Card Hacked
Public Reddit discussion, 2025-09-29. Purposively selected, unverified personal account; no population inference. Unverified reported experience; do not repeat commenters' exploit instructions, bank accusations, amounts or assumed cause.
https://www.reddit.com/r/PersonalFinanceEgypt/comments/1ntct93/

[F02] اشترك في تيلدا بريميوم ولا لا
Public Reddit discussion, 2026-09-12. Purposively selected, unverified personal account; no population inference. Not evidence that listed fees/features remain current. Selected consumer opinions, not representative satisfaction.
https://www.reddit.com/r/PersonalFinanceEgypt/comments/1weaoke/

[F03] حد جرب تطبيق لتتبع المصاريف بشكل أوتوماتيك في مصر؟
Public Reddit discussion, 2026-08-16. Purposively selected, unverified personal account; no population inference. Demand expressed by one self-selected poster; no estimate of market size.
https://www.reddit.com/r/PersonalFinanceEgypt/comments/1vprwoy/

[F04] ازاي اربط تيلدا بinstapay
Public Reddit discussion, 2024-11-03. Purposively selected, unverified personal account; no population inference. Historical feature discussion; not current compatibility documentation.
https://www.reddit.com/r/PersonalFinanceEgypt/comments/1gibrjr/

[F05] اتقفل ... اقبل تحويلات الناس ازاي ؟ Instapay ال
Public Reddit discussion, 2026-10-26. Purposively selected, unverified personal account; no population inference. Cause and account status unverified. Do not repeat advice to evade account controls.
https://www.reddit.com/r/PersonalFinanceEgypt/comments/1xfnf6h/

AI assisted the research, writing, translation, code, artwork and editorial review. This is a working paper, not independently peer reviewed.

softworks labs

Bring your ideas to life.

Work with Softworks. Start with a real problem, build a clear experiment, and measure what changes.

softworks.studio/contact ↗

Interactive edition, data and code: labs.softworks.studio/en/library/egypt-digital-choices
Corrections: send the version, section and source through the studio contact page.