



COGLYTIX RESEARCH

Why Modern Businesses Need Causal AI Decision Intelligence Infrastructure

How evidence, causal reasoning, decision records and accountable action
could work together as a layer between analytics and business execution

WHITEPAPER · REVISED EDITION · SEPTEMBER 2026

DATA



ANALYTICS



DECISION



IMPACT

CONTENTS



ES	Executive summary	03
01	The problem: growth creates noise, and decisions have no home	05
02	Why the existing layers cannot close the gap	09
03	The category: the Causal Decision Layer	11
04	Coglytix Core: the decision intelligence operating model	14
05	The technical foundation: ontology, causal, and generative AI	19
06	Decision briefs, Decision Memory and Declarative Decisioning	23
07	Trust, governance and the deployment model	37
08	Where it pays off	40
09	Why now	48
10	The call to category and the future vision	50
11	Design partner programme	52
AP	About this paper: status, citation and notices	53
GL	Glossary of working terms	54
RF	References	57



ES

Executive summary





The decision layer has no infrastructure

Every modern organisation has invested heavily in two layers of infrastructure. The first moves and stores data. The second turns that data into knowledge and insight through dashboards, models, forecasts, and, increasingly, generative AI summaries. Both layers are mature, well funded, and well understood. Yet the decision itself, the object on which every other investment is meant to pay off, has no infrastructure at all.

The symptom leaders feel first is simpler than any architecture diagram. After every dashboard, forecast, and AI summary, they still do not know what to do. The two existing layers answer what happened, what is happening, and what will likely happen, and then fall silent at the only question that creates value: what action should we take next to produce the outcome we want? Description and prediction are not prescription. Knowing that revenue will likely fall does not reveal which of a dozen possible moves will stop the decline, and no volume of reporting closes that gap.

Decisions are made constantly, but they are made off platform. They sit in dashboards no one acts on consistently, in AI summaries that sound right but can be confidently wrong [1], in spreadsheets, in Slack threads, in email chains, and in data warehouses that hold the evidence but never the verdict. Above all, they sit in human memory. As a result, the decision cannot be extracted, validated, proven, justified, executed, replayed, explained, audited, or tracked. The single most valuable object an organisation produces is also the only one it does not version, store, or measure.

This is the gap. It is not an analytics gap, because analytics is already well served. It is a decision gap, and it requires its own layer of infrastructure.

This paper defines that layer, Decision Intelligence Infrastructure, and the paradigm that makes it real, Declarative Decisioning. It explains why neither analytics nor generative AI can close the gap on its own, why the answer requires a shared ontology beneath a division of labour between causal and generative AI, and why the layer must be built the way the durable infrastructure categories were built: open, declarative, cloud agnostic, free of vendor lock-in, and owned end-to-end by the business, in the lineage of infrastructure-as-code and open data platforms rather than that of closed dashboards and closed decisioning platforms.


The result is a layer designed to do two things at once: to recommend the next best action, and to make that decision testable, explainable, auditable and tracked. It is designed to be reached in two ways, through a declarative code interface for technical teams and a no-code platform for the decision-makers themselves, both resolving to the same underlying specification.

The decision layer begins where analytics ends, helping organisations move from predictions to interventions.



More data than ever. Key business questions still unanswered.

Modern businesses have dashboards, reports, spreadsheets, CRMs, finance tools, marketing platforms, analytics systems, AI assistants and automated summaries. Yet many leaders still struggle with three critical questions:



Analytics tells teams what happened. Dashboards show performance. Reports summarise activity. AI tools generate outputs. But businesses do not win by collecting more information. They win by making better decisions, acting with clarity and learning from outcomes.

This is the gap Coglytix exists to close. Coglytix is building AI-native causal decision intelligence infrastructure for modern businesses, helping teams turn data noise into trusted decisions, accountable action and desired outcomes.

At the heart of this infrastructure sits Coglytix Core, the foundational decision engine beneath every Coglytix product. Core connects business data, context, AI-assisted causal reasoning, decision records, ownership, actions and measurable results in a single accountable loop: from question to decision, from decision to outcome, and from outcome to continuous learning.

THE FOUNDATIONAL LAYER CONNECTS



The purpose of the decision layer is to connect evidence to business impact.



01

The problem: growth creates noise, and decisions have no home






Growth creates noise

In a business's early days, a founder or a small leadership team can hold most of the business decisions in mind. As the business grows, maintaining clarity becomes harder. Data spreads across tools. Teams rely on different metrics. Dashboards and reports multiply. Spreadsheets become information decision hubs. AI generates even more summaries. And yet, important context gets lost across emails, messaging threads, documents and individual memory.

The business has more information, but less clarity.

FIGURE 1 · THE GROWTH NOISE PROBLEM



The problem is not a shortage of data, but the disconnection between data, context, decisions, actions and their outcomes




Decisions have no home

Two failures sit at the centre of how organisations make decisions, and the decision layer exists to solve both. The first is that, even with perfect reporting, nobody can tell what to do. The second is that, on the rare occasions when a sound decision is made, it cannot be preserved.

The question no dashboard answers

Let us take a closer look at a more fundamental decision failure point, aside from the decision preservation problem. Suppose every report were accurate and every forecast were sound. The executive reading them would still not know exactly what to do. Analytics has climbed a familiar ladder, from describing what happened to monitoring what is happening and predicting what is likely to happen. In practice, it stops there, short of prescriptive stage that the standard taxonomy places above prediction [2]. The question that governs the business, namely what action to take to drive the outcome we want, remains unanswered.

FIGURE 2 · THE PRESCRIPTIVE GAP



Analytics answers what happened, what is happening, and what will likely happen. The prescriptive question, what to do next, is the gap the decision layer fills.

This last step is not a harder version of prediction; it is a different kind of question. A forecast that revenue will likely fall is compatible with a dozen possible responses, and nothing in the forecast says which one will work. Choosing the next best action means claiming that the action will cause an improvement, and only a causal model can support that claim [3]. This is why the prescriptive step has resisted the analytics layer for so long: that layer is built on correlation, and prescription requires cause.



Dashboards are valuable. They are not decision systems.

They show what happened but rarely explain what to do next. They display metrics but often do not capture context. They highlight movement but do not automatically create ownership, action or learning. A dashboard may show that revenue slowed. Leadership still needs to understand:

- 01 Why did revenue slow down?
- 02 Which customer segment changed?
- 03 Which channel weakened?
- 04 Was it pricing, churn, conversion, operations or market pressure?
- 05 What options do we have?
- 06 Which action should we take first, and why?
- 07 Who owns it?
- 08 How will we know if it worked?

These questions call for a layer that does not yet exist.

Coglytix is built for the layer that lies between insight and business execution



And then the decision disappears

Consider the lifecycle of an important business decision today, whether it concerns a pricing change, a retention programme, market entry, a hiring strategy or capital allocation. The evidence exists, and it lives in the warehouse. The analysis exists, and it lives in a dashboard or notebook. The discussion exists, and it lives in Slack, Emails and Teams or other meetings tracking and management platforms. The choice gets made. Then the decision itself, including the reasoning, assumptions, expected effect, trade-offs, confidence level and person accountable, simply evaporates. What survives is a number in a deck and a decision in someone's memory.

This produces a set of failures that every operator recognises but few can name. Each maps directly to a capability the that decision layer must restore:

It cannot be **extracted**

The decision and its rationale are scattered across tools and memory, with no canonical record.

It cannot be **validated**

No one checks whether the reasoning was identifiable from the data, or whether the effect was real rather than coincidental.

It cannot be **proven**

When the decision is questioned, there is no counterfactual on record showing what would have happened otherwise.

It cannot be **justified**

The reasoning is reconstructed from memory, after the fact, and selectively.

It cannot be **executed reliably**

The handoff from the moment of choice to a real change in the business is manual and lossy.

It cannot be **replayed**

You cannot re-run the decision against new data to see whether it still holds.

It cannot be **explained**

Stakeholders receive a conclusion rather than a transparent chain of reasoning.

It cannot be **audited**

There is no tamper-evident lineage for regulators, boards, or post-mortems.

It cannot be **tracked**

Predicted impact is never compared systematically to realised impact, so the organisation never learns whether its decisions work.

Software engineering solved the analogous problem with version control. Infrastructure solved it with Infrastructure-as-Code [4]. Decisions remain in the era before both: handcrafted, undocumented, unmeasured, and unrepeatable.



02

Why the existing layers
cannot close the gap






Why the existing layers cannot close the gap

It is tempting to assume that the decision gap is simply the last mile of analytics, or that generative AI will fill it. Both assumptions are wrong and understanding why is at the technical heart of this category. The clearest way to see it is the Ladder of Causation [4], which separates three kinds of questions that a system can answer.

FIGURE 3 · THE LADDER OF CAUSATION



Decisions live on rungs two and three; analytics operates largely, and generative AI wholly, on rung one. Adapted from Pearl and Mackenzie [5]

The ceiling of correlative AI

Nearly all of analytics, and all of generative AI, operates on associations: patterns of the form, what tends to go with what? In the standard hierarchy of reasoning, this is the lowest rung, the rung of seeing, and it answers only the question of what is likely. A decision is not a question about what is likely. It is a question about what happens if we act, which is to say, an intervention. It is a deep and well-established result that we cannot reason validly about an intervention from associations alone [3], [6]. Correlations are confounded; they shift when the world shifts, and they confuse the people who respond to an action with the people who would have acted anyway. An analytics layer, however sophisticated, lives one rung below the place where decisions are made.

The confident wrongness of generative AI

Generative AI is a statistical model of text about the world, rather than a model of the world itself [7]. It is extraordinary at language, synthesis and interface. However, by construction, it is also willing to produce a fluent and persuasive rationale for a conclusion that is causally false [1]. Bolting a large language model onto associative data does not produce decision intelligence. It produces decision theatre: the appearance of rigour without the substance. For a layer whose entire purpose is to prove and justify, this disqualifies generative AI as the foundation, even though, as Section 05 shows, it is indispensable as connective tissue.

The gap, then, is not a tooling gap within the analytics layer. It requires a different kind of reasoning and, therefore, its own layer. Generative AI makes decision-making scalable, while causal AI ensures that the underlying decision logic is sound



When a correlation makes the decision

The gap is not theoretical. In each case below, a correlation that genuinely held in the data was used to support a decision, and that decision was wrong in a way that no amount of predictive accuracy would have exposed [8].

01 A proxy that encoded a bias

A risk-prediction algorithm of a kind applied to roughly 200 million people a year in the United States ranked patients for extra care by predicted health costs, treating spending as a proxy for need. Less care was given to Black patients at the same level of illness, because the algorithm systematically understated how ill they were since the Black patients significantly spent less; Obermeyer and colleagues estimate that removing the proxy would have raised the share of Black patients receiving extra help from 17.7% to 46.5% [9]. The correlation between cost and need was sound in one group and misleading in another, and a predictive model had no way to tell the difference.

02 Risk scores that read a correlate as a cause

Courts in the United States use recidivism scores to inform who is detained. As Sgaier and colleagues observe, the underlying tools rest on statistical correlations rather than causes: low income is correlated with crime without causing it, so people from low-income households can be scored as higher risk and sentenced accordingly [8].

03 Associations that do not survive a test

For years, observational studies linked low vitamin D with diabetes, hypertension, cardiovascular disease and cancer. Randomised trials then found that supplementation did not lower the incidence of invasive cancer or cardiovascular events [10], nor the risk of type 2 diabetes [11]. The association was genuine; the causal claim that an intervention would have needed was not.


04 Why experiments alone cannot close the gap

Randomised trials remain the gold standard, and the 2019 Nobel Prize in economics recognised their use against global poverty [12]. But they are costly and slow, can only test one or a few bundled interventions at a time, and estimate the effect on a population rather than on the case at hand [8]. Causal methods applied to data the organisation already holds are how that logic reaches everyday decisions.



03

The category: the
Causal Decision
Layer






The Causal Decision Layer

Once the missing layer is named, the shape of the enterprise stack becomes obvious. The diagram below places it where it belongs, between the layers that produce knowledge and the layer that produces impact.

FIGURE 4 · THE ENTERPRISE STACK



The enterprise stack, with the decision layer in its place.


The lower two layers turn data into knowledge. The decision layer turns knowledge into justified action. It is designed to help narrow what Pfeffer and Sutton called the knowing-doing gap [13], carrying the familiar progression from data to information, knowledge, understanding and wisdom [14] on to choosing the next action, with the reasoning and level of confidence preserved as a first-class, durable artefact. It is the layer where an insight stops being a slide and becomes an accountable, auditable intervention with a measured outcome.

This is a feature of the decision layer, with its own primitives, lifecycle and infrastructure.



From analytics to decision intelligence

FIGURE 5 · FROM ANALYTICS TO DECISION INTELLIGENCE



Coglytix is built for the layer between insight and impact. The decision layer is where insight becomes accountable action


Modern businesses do not need more noise. They need decision intelligence infrastructure.

Decision intelligence infrastructure connects five things that are usually separated:

- 1 Business data
- 2 Business context
- 3 AI-assisted causal analysis
- 4 Human decision making
- 5 Action & impact tracking

This creates a continuous operating loop. Data becomes insight. Insight becomes decision. Decision becomes action. Action becomes impact. Impact becomes learning.

FIGURE 6 · THE COGLYTIX VALUE CHAIN




Learning feeds back into business context



04

Coglytix Core: the
decision intelligence
operating model






Five critical questions, one foundational layer.

Coglytix Core is the foundational layer of decision intelligence, organised around five questions that every consequential decision raises.

FIGURE 7 · THE FIVE QUESTIONS COGLYTIX HELPS ANSWER



Coglytix Core is neither a dashboard, nor a reporting tool, nor an AI chatbot.

Coglytix is the decision intelligence infrastructure for modern businesses.




Four primary capabilities

FIGURE 8 · THE COGLYTIX FRAMEWORK



WHAT THE FRAMEWORK PRODUCES





Seven steps. One continuous loop.

FIGURE 9 · THE DECISION INTELLIGENCE LOOP



learning feeds the next signal

This loop turns decision making into an organisational capability.



Six architectural layers

FIGURE 10 · COGLYTIX CONCEPTUAL ARCHITECTURE






A decision layer across your existing tools

The system is designed to act as a decision layer across existing business tools, not as a forced replacement for them.

FIGURE 11 · COGLYTIX AS THE DECISION LAYER




Coglytix draws on the systems teams already use and adds the decision layer above them, providing the needed clarity rather than another silo.



05

The technical foundation:
ontology, causal, and
generative AI





The technical foundation: ontology, causal, and generative AI

The decision layer rests on a shared ontology, with a clear division of labour between the two AI paradigms above it, each performing the function for which it is best suited.

01 A shared ontology is the foundation

Before any causal claim can be trusted, the entities it concerns must mean one thing across the business. If customer, churn, and margin are defined differently in the warehouse, the CRM, and the boardroom, then a causal estimate built on them is quietly wrong no matter how sound the mathematics. An ontology is a formal, shared model of the organisation's entities, relationships, and metrics, defined once. It is what makes a causal graph valid, what lets a decision written once be understood everywhere, and what gives generative extraction a schema to map language onto rather than a blank page to invent on. The ontology is not a feature bolted on top; it is the ground the rest of the layer stands on.

02 Causal AI is the decision engine

It models the mechanisms that generate the data rather than the patterns on the surface, which lets it answer the two questions a decision poses: what happens if we act, and what would have happened otherwise. It supplies the effect and, critically, the confidence. Its toolkit ranges from randomised experiments, which are the gold standard [16], to quasi-experimental methods for when randomisation is impossible, such as difference-in-differences, regression discontinuity, instrumental variables [17], and synthetic control [18]. It extends further to heterogeneous-effect models that identify who responds [19], to double and debiased machine learning [15], and to sensitivity analysis that bounds the risk of unmeasured confounding [20]. Because a recommendation is a claim that an action will cause an outcome, only a model of cause can produce a trustworthy next best action, whereas correlative prescription merely repeats whatever happened to coincide with success in the past.

03 Generative AI is the interface and the connective tissue


It does the work no causal engine can. It reads the unstructured chaos where decisions are trapped, including documents, dashboards, Slack, and email, and it extracts the candidate decisions, the hypotheses, and the variables in play. It renders causal results into language an executive can act on, drafts the justification and the memo, and orchestrates the workflow. It is the layer's senses and its voice, but not its judgment.

Together they form a closed loop beneath the operating loop of Section 04. Open-source libraries now make its causal steps deployable in practice [21], [22].



The closed decision loop

FIGURE 12 · THE CLOSED DECISION LOOP



Causal AI supplies the judgement; generative AI supplies the interface.

- 1 Extract (generative).** Recover candidate decisions and variables from unstructured sources. This is where decisions stop living only in human memory.
- 2 Structure (causal).** Build the causal model and determine what is identifiable from the available data and the stated assumptions.
- 3 Estimate (causal).** Quantify intervention effects with explicit uncertainty, using experiments or the appropriate quasi-experimental method.
- 4 Simulate (causal).** Run interventions forward to ask what if we do this, and counterfactuals backward to ask what would have happened [23].
- 5 Optimise (decision theory).** Select the next best action, the one that maximises the objective under the causal model, the constraints, and the organisation's risk tolerance, and quantify the value of gathering more information before committing [24]. This step is the direct answer to the prescriptive question that analytics leaves open.
- 6 Explain and record (generative).** Write the justification and commit the decision as a versioned artifact.
- 7 Track and learn.** Compare predicted effect to realised effect, so that the drift recalibrates the model and feeds the audit trail.

A NOTE ON CONFIDENCE, STATED PLAINLY

High and defensible confidence is a product of method rigour, namely identification, experimentation, and uncertainty quantification, rather than of any model's fluency. In some domains, irreducible uncertainty means a high stated confidence is simply unattainable, and there the value of the layer is to quantify and bound the uncertainty honestly rather than to manufacture false certainty. A decision layer that knows what it does not know is worth more than one that is confidently wrong.



Two paths to a causal answer

Causal AI is not a single technique. Two long-established families of methods answer interventional questions using data that an organisation already holds. They differ in the number of causes they can evaluate at once [8].

01 The potential-outcomes route

Compare what happened to an exposed unit with what would have happened to it had it not been exposed. Where no control group exists, comparable unexposed units are matched to construct one, effectively reverse-engineering the trial that was never conducted; propensity-score matching, introduced by Rosenbaum and Rubin is among the most widely used methods [25],[8]. Its strength lies in the rigour it brings to a question that can be stated in advance. Its limitation is scope, since each analysis evaluates a single intervention named beforehand [8].

02 The causal-graph route

Structural models and causal Bayesian networks [3], [26] map the web of pathways to an outcome rather than a single pair of variables. A structural model tests the relationships an analyst specifies; a Bayesian network can estimate relationships across all the variables in a dataset, which makes it a discovery method as well as a test and lets many interventions be simulated together [8]. Its cost is data: whatever part of the relevant world is missing from the dataset is missing from the graph.

03 Neither is a black box, and neither runs without people

In both routes the relationships between variables and outcome stay visible and can be checked by those who must act on them, a transparency that deep-learning predictors do not provide [8]. Expert knowledge enters explicitly: practitioners constrain the model, name confounders the data would otherwise omit, and interpret results that run against received wisdom [8].

The decision layer needs both: potential-outcomes estimation when the intervention is already named, and graph-based discovery when the drivers are still in question.



AI supports judgement. It does not replace it.

Coglytix uses causal AI as a support layer for better decision-making. The goal is not to replace leadership judgement, but to improve clarity, speed, structure and learning.

Generative AI learns correlations, and scale makes it more fluent [27]. Causal AI models cause and effect, which is what any claim about the result of an action requires [3].





FIGURE 13 · HUMAN-CENTRED AI DECISION SUPPORT



Coglytix's approach to AI is grounded in causal reasoning, enabling decision support that is explainable, accountable, and measurable.



06


Decision briefs,
Decision Memory and
Declarative Decisioning





From scattered signals to a structured decision object

FIGURE 14 · ANATOMY OF A DECISION BRIEF



WITHOUT A BRIEF

"We need to improve performance."

Vague conversations, unclear ownership, no follow-through.

WITH A BRIEF

A brief records what has changed and why it matters, the options open to us, the course chosen, who is accountable for it and how success will be judged. The pages that follow complete it for eight industries; every organisation and figure in them is invented.



The method behind every brief

The eight briefs that follow are worked examples, not case studies. They follow the same procedure but do not assume the same conclusion. Each brief shows how every claim was derived.

- 01 State who is eligible for the decision, and who has deliberately been excluded from it.
- 02 Describe the intervention closely enough that another team could administer it unaided.
- 03 Define the comparator. In a regulated process, the alternative is rarely to do nothing.
- 04 Establish the outcome, and the period over which it is measured, before any of the work begins.
- 05 Specify the estimand: which incremental outcome, for whom, against what, and over how long.
- 06 Apply the eligibility, policy and safety rules, and record which cases they exclude.
- 07 Value each option by a stated formula, giving a range and a downside case wherever the data allow.
- 08 Record the decision, the reasoning behind it and the action that was taken.
- 09 Evaluate the outcome using the method selected in advance, not one chosen to favour the result.
- 10 Publish the outcome review separately after sufficient evidence has accumulated.

ON THE STATUS OF THESE BRIEFS

The organisations and figures used in these briefs are illustrative.

They are included to show how the approach works, not to represent actual organisations or results.

The structure is what matters here, rather than the specific numbers.

Each recommendation is a hypothesis to be tested, not a conclusion already reached.

In practice, every recommendation would remain subject to eligibility criteria, organisational policies, legal and regulatory requirements, data-quality standards and human approval.

These examples should not be interpreted as claims about Coglytix's performance.



Decision brief: fraud intervention

In payments, a transaction has been flagged as suspicious. The question is not whether the model is right, but which of the permitted responses is worth making, and what it costs the customer who was never a fraudster.

DECISION BRIEF

THE SITUATION

Issue	A fixed score threshold sends every flagged payment to decline or review, whatever the case.
Signal	3,100 of 190,000 payments flagged each day; 1.4% falsely declined; review capacity 900 a day.
Evidence	E2. A year of payment records shows which payments look risky, not which response works.


THE CHOICE

Options	Approve, step-up authentication, customer confirmation, short hold, review, decline or restrict.
Comparator	The current compliant fraud-control policy.
Estimand	Net fraud-control value per 1,000 eligible flagged payments, once fraud labels mature.
Causal design	Responses randomised within risk-score bands among eligible payments.
Expected value	Illustrative: central £46,000; range £28,000 to £61,000; downside a loss of £19,000.
Main risk	Step-up authentication may drive away customers who were never fraudsters.
Constraint	Strong customer authentication applies [28], [29]; refusal must be justified and notified, though a payment may be delayed on suspicion of fraud [30], [31]; scam reimbursement applies [32].
Recommendation	Pilot risk-sensitive responses within the space left by mandatory controls. Confidence: moderate.

THE COMMITMENT

Owner	Head of Fraud Strategy.
Success measure	Incremental net value per 1,000 payments; guardrails on false declines, complaints and fairness.
Review	Daily safety checks; full evaluation once fraud and reimbursement outcomes have matured.

CAUSAL STRUCTURE



Randomising within score bands removes the path from risk to response, so any difference in outcome can be attributed to the response itself. The design assumes that one payment's treatment does not alter another's outcome.



Decision brief: arrears intervention

In household energy, a customer has fallen into arrears. What matters is not whether contact produces a payment, but whether the arrangement that follows is one the household can sustain.

DECISION BRIEF

THE SITUATION

Issue	One collections sequence treats temporary difficulty and persistent hardship alike.
Signal	7,200 accounts enter arrears each month; 27% of new arrangements break within 90 days.
Evidence	E2. A year of billing records shows who repays, not what makes repayment hold.


THE CHOICE

Options	Reminder, affordability review, tailored plan, specialist referral or billing check.
Comparator	The current compliant collections sequence.
Estimand	Share of eligible customers holding an affordable arrangement at 90 days.
Causal design	Eligible accounts randomly allocated between compliant pathways.
Expected value	Illustrative: central £24,000; range £16,000 to £33,000; a loss of £12,000 if arrangements fail.
Main risk	An arrangement the household cannot afford deepens the harm it was meant to relieve.
Constraint	Suppliers must take account of ability to pay [33]; vulnerable and disputed cases go to specialist teams and are never placed in a lower-support group.
Recommendation	Compare compliant early engagement. Confidence: moderate on segmentation, low on effect.

THE COMMITMENT

Owner	Director of Customer Operations.
Success measure	Affordable arrangements in place at 90 days; guardrails on complaints and repeat arrears.
Review	Ninety days after the final enrolment, against the randomised comparison group.

CAUSAL STRUCTURE



Vulnerable and disputed cases are excluded before allocation. Among the remaining accounts, random allocation ensures that payment history cannot influence which pathway a customer receives.



Decision brief: production reliability

On an offshore installation, a compressor is deteriorating. Monitoring tells us that failure has become more likely; it does not tell us whether intervening now is worth the production it would cost.

DECISION BRIEF

THE SITUATION

Issue	A critical compressor is deteriorating; the planned maintenance window closes in 12 days.
Signal	34% failure risk in 30 days; £145,000 a day at risk; repair after failure up to £480,000.
Evidence	E2. Sensor and service history predict failure, not the effect of any treatment.


THE CHOICE

Options	Continue, monitor closely, reduce load, maintain in the window, or shut down now.
Comparator	The existing monitoring and maintenance plan.
Estimand	Reduction in unplanned downtime over the 90 days after the window.
Causal design	Matched fleet units compared on failure and downtime, using survival analysis.
Expected value	Illustrative, versus continuing: maintain £285,000; shut down £170,000; reduce load £125,000.
Main risk	The diagnosis may be wrong, and the work itself may introduce a new fault.
Constraint	Safety case [34], pressure-system maintenance duty [35] and a written permit to work [36].
Recommendation	Maintain once inspection confirms the fault; reduce load meanwhile. Confidence: moderate.

THE COMMITMENT

Owner	Asset Operations Manager.
Success measure	Unplanned downtime against matched units; guardrails on safety and environmental incidents.
Review	At 14 and 90 days after maintenance, and at the next major maintenance cycle.

CAUSAL STRUCTURE



A single asset cannot be randomised, so comparable units across the fleet are matched on the factors that drive both maintenance timing and failure. The estimate holds only if no unmeasured factor does the same.



Decision brief: early delinquency

In consumer credit, a borrower has missed a single payment. Most such borrowers recover unaided, which is precisely why the effect of intervening is so hard to read from the record of past interventions.

DECISION BRIEF

THE SITUATION

Issue	Every first missed payment receives the same contact sequence, whatever its cause.
Signal	2,100 first missed payments a month; 14% go on to reach 60 days past due.
Evidence	E2. Payment and income records predict who falls further behind, not which support prevents it.


THE CHOICE

Options	Digital reminder, payment-date change where permitted, short arrangement, forbearance assessment or human contact.
Comparator	The lender's existing compliant process.
Estimand	Reduction in progression to persistent arrears among eligible accounts at 90 days.
Causal design	Eligible accounts randomised between compliant pathways; no one is denied due forbearance.
Expected value	Illustrative: central £27,000; range £18,000 to £37,000; downside a loss of £14,000.
Main risk	Support may postpone arrears rather than prevent them.
Constraint	Forbearance and due consideration for every customer in arrears [37]; affordability allows for priority debts and essential spending.
Recommendation	Compare compliant pathways, keeping vulnerable cases with staff. Confidence: moderate.

THE COMMITMENT

Owner	Head of Credit Operations.
Success measure	Fewer accounts reaching persistent arrears; guardrails on sustainability and fair outcomes.
Review	Ninety days after the final intervention, reported by subgroup.

CAUSAL STRUCTURE



Randomisation makes borrowers in each pathway alike on average, so the difference in arrears reflects the support itself rather than the kind of borrower who happened to receive it.



Decision brief: churn intervention

In telecommunications, a contract is approaching its end. The difficulty is that the customers most likely to leave and those most responsive to an offer are rarely the same people.

DECISION BRIEF

THE SITUATION

Issue	Discounts go to the highest churn scores, many of whom would have stayed anyway.
Signal	11,500 of 48,000 renewing customers scored high risk; 71% of discounted customers renew.
Evidence	E2. Churn scores predict who leaves; they say nothing about whom an offer persuades.


THE CHOICE

Options	Renewal notice only, service review, plan redesign, targeted discount or human contact.
Comparator	The standard renewal communication.
Estimand	Incremental informed renewal within 30 days of contract expiry, by segment.
Causal design	Customers randomised across interventions, so uplift is measured rather than inferred.
Expected value	Illustrative: central £22,700; range £14,300 to £31,100; downside a loss of £13,000.
Main risk	Retention bought with discounts can conceal a service problem.
Constraint	No barriers to switching [38]; Ofcom fined one provider £28m for such barriers in 2026 [39].
Recommendation	Target discounts by estimated uplift and value, not by churn risk alone. Confidence: moderate.

THE COMMITMENT

Owner	Customer Retention Director.
Success measure	Incremental retained contribution; guardrails on ease of cancellation and complaints.
Review	Ninety days after the final renewal deadline in the pilot.

CAUSAL STRUCTURE



Randomisation separates the customers an offer persuades from those who would have renewed regardless, a distinction that no churn score, however accurate, can draw.



Decision brief: home insurance renewal

In retail home insurance, a policy is approaching renewal. Retention is easy to buy and easy to mistake for a good outcome, particularly where the cover no longer suits the household.

DECISION BRIEF

THE SITUATION

Issue	Lapse is rising among profitable home policies, and retention effort follows predicted lapse risk.
Signal	30,000 policies renewing; renewal rate 76%; 6,400 retention offers at £38 each.
Evidence	E2. Renewal tracks price, claims and service, which does not show what prevents lapse.


THE CHOICE

Options	Notice only, needs review, service recovery, payment plan or a rule-compliant offer.
Comparator	The standard renewal process.
Estimand	Incremental suitable renewal within 30 days of expiry.
Causal design	Eligible policies randomised across compliant interventions; price never varies outside the rules.
Expected value	Illustrative: central £7,400; range £4,150 to £10,675; downside a loss of £6,000.
Main risk	Retention may keep a household on cover that no longer suits it.
Constraint	Renewal price may not exceed the equivalent new-business price [40]; fair value applies.
Recommendation	Test non-price interventions before any price change. Confidence: moderate.

THE COMMITMENT

Owner	Director of Customer Retention.
Success measure	Incremental suitable renewals; guardrails on complaints and vulnerable customers.
Review	Ninety days after the final renewal date in the pilot.

CAUSAL STRUCTURE



Random assignment balances claims history and tenure across groups, so a difference in suitable renewals can be credited to the intervention rather than to the policyholder.



Decision brief: route profitability

In contract logistics, fourteen of thirty-eight routes lose money. The costs are known with reasonable precision; what each remedy would do to them is not.

DECISION BRIEF

THE SITUATION

Issue	A £2.8 million contract loses money on 14 of its 38 routes.
Signal	Contribution negative £145,000 a year; failed first delivery 9.5%; average wait 47 minutes.
Evidence	E2. Cost analysis locates the losses; simulation predicts savings that are not yet causal.


THE CHOICE

Options	Continue, reprice, redesign routes, change service levels, or a combination.
Comparator	The current contract and route configuration.
Estimand	Incremental contribution per completed delivery over twelve weeks.
Causal design	Route clusters phased in at random and compared by difference-in-differences.
Expected value	Illustrative: central £123,000; range £97,000 to £149,000; a loss of £28,000 if volumes fall.
Main risk	Repricing may lose the customer; consolidation may lengthen delivery times.
Constraint	Contractual notice, delivery-window commitments, service penalties and driver capacity.
Recommendation	Redesign routes in clusters before any contract-wide repricing. Confidence: moderate on cost, low on network effects.

THE COMMITMENT

Owner	Operations Director.
Success measure	Contribution per delivery against comparison clusters, adjusted for fuel, volume and season.
Review	Fourteen weeks after the pilot begins, with spillovers between clusters checked.

CAUSAL STRUCTURE



Comparing the change over time against untreated clusters removes fixed differences between routes, provided both groups would otherwise have followed parallel trends, which the pre-period allows us to check.



Decision brief: inventory and working capital

In physical commodity trading, inventory sits at one port while prices are higher at another. The market itself cannot be intervened upon, so the causal question concerns the desk's own allocation policy.

DECISION BRIEF

THE SITUATION

Issue	Storage and financing costs are rising on 20,000 tonnes held at Port A.
Signal	£8.4 million at stake; Port B premium £18 a tonne against £10 transport; transit 12 to 18 days.
Evidence	E0 to E2. Holdings and obligations are observed facts; prices and spreads are forecasts.


THE CHOICE

Options	Hold, sell now, move all, move part, sell forward before transport, or reallocate.
Comparator	Holding for 30 days, a benchmark fixed before the outcome is known.
Estimand	Risk-adjusted value of each allocation against that benchmark, at one valuation timestamp.
Causal design	Allocation rules alternated across comparable cargoes over time, in a switchback design.
Expected value	Illustrative, against holding: move all £118,000; move half £106,000; forward sale £101,000. Acquisition cost is sunk.
Main risk	The spread may close before the cargo arrives.
Constraint	Delivery commitments, credit and risk limits, hedging policy, duties and sanctions screening.
Recommendation	Move half under a forward sale at Port B; moving all scores higher but carries the full spread risk. Confidence: low to moderate.

THE COMMITMENT

Owner	Head of Physical Trading.
Success measure	Realised value against the benchmark, split into market and execution effects.
Review	At delivery, and 30 days after the decision.


CAUSAL STRUCTURE



Prices lie outside the desk's control, so the switchback design isolates the policy: because the alternation is fixed in advance, market conditions cannot determine which rule a cargo receives.




In most businesses, decisions disappear



When this happens, organisations lose learning. Coglytix introduces decision memory, capturing what was decided, why it was decided, who made the decision, what evidence was used, what action was agreed, who owned it, what result followed and what the business learned.

FIGURE 15 · DECISION MEMORY



Decision memory turns each business decision into a learning asset.




The paradigm: Declarative Decisioning

For the decision layer to function as infrastructure rather than as an application, it requires its own operating paradigm. That paradigm is Declarative Decisioning, in which decisions are expressed as declarative specifications and managed through a lifecycle of Plan, Apply and Track.

The Decision Specification. A decision is declared, not clicked together. The specification states the objective, causal assumptions expressed as a structural model, data sources, constraints, risk tolerance and decision policy. Because it is written declaratively, the decision becomes reviewable, comparable across versions and version-controlled. This represents the same leap that Infrastructure as Code brought to servers [4].

The Lifecycle. The Plan step computes the predicted effect and its level of confidence before anything is committed. It also shows the reasoning, much as a dry run reveals the likely consequences before action is taken. The Apply step executes the decision through the systems of record. The Track step measures the realised effect and reconciles it with the prediction. The difference between the planned and realised impact provides both a continuous calibration signal and standing proof of value.

FIGURE 16 · THE DECLARATIVE DECISIONING LIFECYCLE



The Declarative Decisioning lifecycle, with every step captured in the signed Decision Memory.

THE SIGNED DECISION MEMORY · THE MOAT

It is designed as a cryptographically signed, hash-chained, append-only record of every decision, capturing its inputs, assumptions, predicted effect, confidence, accountable owner, and realised outcome. Because each entry is signed and chained to the last, the record is tamper-evident and independently verifiable [41]: any decision can be replayed exactly as it was made and proven to a regulator or a board to have been made that way. It is the single artefact that addresses most of the failures in Section 01, because it brings to decisions the discipline that version control brought to source code, with cryptographic assurance of integrity added.. Extraction, validation, proof, justification, replay, audit, and tracking stop being aspirations and become operations against the signed Decision Memory.




Two interfaces to one engine

The paradigm does not dictate a single way in. The same engine, Coglytix Core, the foundational layer described in Section 04, along with the same lifecycle and signed Decision Memory, can be accessed through two interfaces, each suited to a different audience.

Declarative Decisioning decisions authored as declarative specifications, reviewed and versioned like any other code. It is the natural interface for data and platform engineers. It serves as the source of truth.

Decision Studio It is a no-code visual interface delivered as a cloud-based solution and designed for the people who act on decisions, namely executives and leaders across finance, operations and people functions. They define the objective, assumptions, constraints and policy through a guided interface rather than a text file. Crucially, every action within the platform compiles into the same underlying Decision Specification. The platform introduces no separate logic of its own. It provides a view of the same specification rather than creating separate functionality.

FIGURE 17 · TWO INTERFACES TO ONE ENGINE



Because both interfaces produce the same artefact, a decision created in the platform can be exported, read and edited as code. Likewise, a decision authored as code can be opened and rendered in the platform. Nothing is trapped in the user interface.

This matters most to the people who experience the prescriptive gap directly. The leaders responsible for choosing the next best action are rarely the same people who write specifications by hand. The platform makes the decision layer accessible to them without sacrificing the rigour, portability and auditability guaranteed by the code interface.



The decision specification

The specification is the artefact on which every other step operates. The example below is illustrative rather than representative of the product's syntax. What matters is that every element of a decision is declared once in text that can be reviewed, compared across versions and version-controlled.

decision-spec.yaml

```
decision: winter_retention_offer
objective:
  maximise: incremental_gross_margin
  horizon: 90d
assumptions:
  graph: retention_v3          # structural model
  treatment: discount_offer
  outcome: renewal
  confounders: [tenure, plan, support_contacts, usage_trend]
data:
  sources: [warehouse.customers, warehouse.billing]
  ontology: core@2026-09      # shared definitions
constraints:
  budget_max: 250000
  exclude: [accounts_in_dispute]
  human_review: required
risk_tolerance:
  min_confidence: 0.85
policy:
  target: customers_with_positive_uplift
  review_after: 30d
owner: director_of_retention
```

01 plan

Returns the predicted effect with its interval, the assumptions the estimate depends on, and what the decision costs if an assumption fails. Nothing has been committed yet, and the reasoning is reviewable before it is.

02 apply

Executes the decision into the systems of record and assigns the owner named in the specification, so the handoff from choice to change is not manual.

03 track

Compares realised effect with predicted effect and files the drift against this version of the specification, which is what makes the next version better informed than this one.



07

Trust, governance and the deployment model






Decision intelligence must be trusted

An organisation cannot rely on recommendations if it does not understand the evidence, assumptions and context behind them. Coglytix is guided by five trust principles of its own. Two of these principles: explainability and accountability, correspond to characteristics of trustworthy AI in the NIST AI Risk Management Framework [42]; the remainder reflect what decision records require.

FIGURE 18 · TRUST MODEL






The deployment model: open, sovereign, customer owned

This is where the decision layer is most open to honest scrutiny and where its strongest claim lies. There are already some players in this space. Decision intelligence is not a new idea: the discipline was named and developed by Lorien Pratt and Mark Zangari, among others [43]. A generation of decision intelligence platforms already exists, and some combine causal and generative AI. In our assessment, most are proprietary, vendor-hosted services. The customer licenses the vendor's cloud, does not own the underlying infrastructure and, in regulated industries, may be unable to keep data in-house.

That is the opportunity. The decision layer does not claim to have invented decision intelligence. It is an open, customer-owned alternative to the closed platforms already serving this market. It follows the declarative, provider-based model of Infrastructure as Code tools such as Terraform and OpenTofu rather than the closed Software as a Service model.

FIGURE 19 · CONTROL PLANE AND DATA PLANE



Control-plane and data-plane separation keeps data and compute inside the customer's boundary.

SOVEREIGN DEPLOYMENT FROM ONE CODEBASE

The control plane orchestrates, while the data plane, comprising the compute and data, runs within the customer's own environment. This may be a cloud account, a virtual private cloud or a fully on-premises deployment. The same codebase is designed to support all three, with no separate or weaker on-premises edition.

Data never crosses the customer's boundary. The promise that customers own the infrastructure end to end is therefore not a marketing statement but an architectural commitment. This makes the layer suitable for highly regulated industries, where sending data to a vendor-hosted service is often restricted.



Built the way durable infrastructure is built

01 Regulatory-native from day one

Because every decision is captured in the signed Decision Memory, the layer is designed around the EU AI Act [44] and the NIST AI Risk Management Framework [42] from the outset rather than retrofitted to meet their requirements. The signed, replayable record is designed to support the Act's record-keeping (Article 12), human-oversight (Article 14) and deployer (Article 26) obligations for high-risk AI: what was decided, on what basis, with what predicted effect, and what actually happened. Compliance evidence therefore becomes a by-product of the architecture.

02 Cloud agnostic and declarative

The same declarative specification deploys against any cloud and any warehouse, with no rewrite required to move from one provider to another.

03 A provider model, and no lock-in

Following the pattern of the Terraform Registry and dbt packages, the platform offers pluggable connectors to data platforms and communication tools, pluggable causal engines, and the freedom to bring your own language model. The decision specifications and the signed Decision Memory use open formats, so the artefacts remain portable. The customer can leave at any time, and the fact that they choose not to is the only acceptable kind of lock-in.

04 A marketplace of decision packs

Prebuilt, parameterised decision modules address each use case, such as a pricing-elasticity pack, a churn-uplift pack, or a causal-attribution pack, distributed in the manner of infrastructure modules or apps. This is the distribution engine and the path to an ecosystem in which third parties extend the platform.



08

Where it pays off





Where causality changes the decision

The same loop applies across the enterprise. The causal question changes from one domain to the next, but the pattern does not.

REVENUE

Does this discount cause the purchase, or merely target people who would have bought anyway?

Method: Uplift modelling [45], incrementality tests, causal attribution

Flagship: Incremental-uplift targeting: spend only where the action causes revenue, with counterfactual proof on record

CUSTOMER

Which retention actions change what customers do, and for which customers do they work?

Method: Retention uplift modelling [45], heterogeneous treatment effects [19]

Flagship: Churn-uplift targeting: retention offers go only to customers they persuade to stay, with realised effects tracked by segment

PEOPLE

What causes attrition and performance, and which pathways are legitimate rather than illegitimate?

Method: Quasi-experiments, counterfactual fairness analysis [46]

Flagship: Causal attrition drivers with intervention effects by segment, kept as human-reviewed decision support

RISK

What is the true root cause, and what happens under stress conditions not yet observed?

Method: Root-cause structural models, counterfactual stress scenarios

Flagship: Counterfactual scenario simulation for early warning and capital at risk

FINANCE

What is the incremental return on spend, and will the forecast survive a change of regime?

Method: Incrementality testing, causal forecasting of mechanisms

Flagship: Causal driver forecasting that is more robust to distribution shift because it models mechanisms [47]

STRATEGY

What would happen if we made this irreversible bet?

Method: Interventional simulation, counterfactual analysis of timing

Flagship: Big-bet simulation for market entry and M&A, modelled as interventions with bounded confidence

OPERATIONS

Does this process change cause the throughput gain we expect?

Method: Causal effect estimation, causal digital twins

Flagship: Causal digital twin that answers change this lever, get this effect, this sure

A caution in the People domain. Causal methods are the responsible choice precisely because they can separate a legitimate pathway from an illegitimate one [48], [49], but effect estimates that touch sensitive or protected attributes belong in human-reviewed decision support and never in automation.



Where we are proving it first

The preceding page sets out where the method applies. This one is deliberately narrower: the four economic outcomes on which our first deployments with design partners are focused.

- 01** Cash flow and cash recovery. Which action recovers money that would otherwise be written off, and how much of that recovery would have occurred in any case? Measured as incremental cash recovered per eligible account, net of the cost of acting.
- 02** Customer retention and revenue protection. Which customers remain because of what was done, rather than because they were never going to leave? Measured as incremental retained contribution per eligible customer, net of the offer.
- 03** Profit intelligence. Which products, routes, contracts and customers cover their own cost once the true cost to serve is carried? Measured as incremental contribution per unit of delivered work, against the current operating plan.
- 04** Growth efficiency. Which expenditure causes growth, and which merely accompanies it? Measured as incremental outcome per pound deployed, against a comparator fixed before the result is known.

ON SCOPE, AND ITS LIMITS

These four constitute the initial deployment scope; they are not the boundary of the platform.

They were chosen because the money is attributable and the comparator reasonably clear.

A deployment is counted as proven only where the comparator was fixed in advance, the effect estimated by a pre-specified method, and the realised value reconciled against that estimate.

The decision layer is designed to be domain-neutral, able to hold any decision with options, evidence and an owner.

Beyond these four the method is unchanged; only the quantity being measured differs.

Breadth is a property of the method; proof is a property of a deployment. This page concerns the latter, and the two that follow show the record such a deployment leaves behind.



Recovered cash and retained customers

Two of the four outcomes, recorded as they would appear in Coglytix Growth, the first product built on Core, for UK business growth-stage and scaling businesses.

DECISION RECORD · COGLYTIX GROWTH

Recover overdue invoices without straining the relationship

Signal	Days sales outstanding rising across a book of 1,800 trade accounts
Why	Evidence points to reminder timing and payment friction, not an intent to delay
Options	4 collection pathways ranked by expected incremental economic value
Decided	Move eligible accounts to an earlier, lighter reminder carrying a one-click payment link
Comparator	Eligible accounts randomly held on the existing reminder sequence
Outcome	A higher share of invoices settled within 30 days after the change
Incremental impact	Estimated uplift versus comparable untreated accounts: +5.2 percentage points
Economic value	Estimated incremental cash recovered: £41,300 in one quarter
Evidence	Randomised holdout · E4 · 88% confidence
Learning	Earlier and lighter contact worked for this segment; the harder sequence did not improve on it. The intervention, context, outcome and evidence are kept in the signed Decision Memory to inform future decisions.

DECISION RECORD · COGLYTIX GROWTH

Reduce churn in mid-tier accounts

Signal	Repeat-purchase rate declining across a mid-tier customer segment
Why	Evidence points to onboarding gaps as the most likely actionable driver, rather than price
Options	3 interventions ranked by expected incremental economic value
Decided	Introduce guided onboarding for eligible new mid-tier accounts
Comparator	Accounts opened just before rollout, by regression discontinuity; balance and density tests passed
Outcome	Repeat purchase improved after the intervention
Incremental impact	Estimated uplift versus comparable untreated accounts: +7.4 percentage points
Economic value	Estimated incremental contribution protected: £18,400
Evidence	Quasi-experimental estimate · E4 · 91% confidence
Learning	Guided onboarding increased repeat purchase for this segment. The intervention, context, outcome and evidence are kept in the signed Decision Memory to inform future decisions.

EVIDENCE LEVELS

E0 denotes an observed fact; E1 a measured outcome; E2 a predictive association; E3 a quasi-experimental estimate whose identifying assumptions cannot be fully tested; and E4 an experimental estimate, or a quasi-experimental one whose testable assumptions have been checked and not rejected. The level travels with the number, so that a reader may judge how much weight it carries before acting upon it.



Protected margin and efficient growth

The remaining two. A brief is written before the decision is taken; a record is what survives once the result is in and the evidence has been reconciled against what was predicted.

DECISION RECORD · COGLYTIX GROWTH

Stop losing money on small, low-density orders

Signal	Contribution falling in the smallest order band while its volume rises
Why	Evidence points to cost to serve and delivery frequency, not list price
Options	3 changes ranked by expected incremental economic value
Decided	Set a minimum order value and a consolidated delivery day for eligible accounts
Comparator	Matched accounts in the same order band outside the treated delivery areas
Outcome	Contribution per order improved in the treated band
Incremental impact	Estimated uplift versus comparable untreated accounts: +£2.80 per order
Economic value	Estimated incremental contribution: £26,700 annualised
Evidence	Matched difference-in-differences · E3 · 84% confidence
Learning	Cost to serve, driven by small orders and frequent deliveries, carried the loss; price was not the cause. The intervention, context, outcome and evidence are kept in the signed Decision Memory to inform future decisions.

DECISION RECORD · COGLYTIX GROWTH

Find the paid spend that drives new business

Signal	Acquisition cost rising while qualified enquiries stay flat
Why	Evidence points to overlap between paid and organic demand, not creative quality
Options	3 budget allocations ranked by expected incremental economic value
Decided	Pause paid search in two regions for six weeks and hold the rest unchanged
Comparator	Regions assigned at random to continue paid search unchanged
Outcome	Enquiry volume in the paused regions held close to its previous level
Incremental impact	Enquiries genuinely caused by paid search: 18% of those previously credited to it
Economic value	Estimated spend released, net of the enquiries lost: £14,900 a quarter
Evidence	Geographic holdout test · E4 · 90% confidence
Learning	Most enquiries attributed to paid search in those regions would have arrived anyway. The intervention, context, outcome and evidence are kept in the signed Decision Memory to inform future decisions.

ON THE STATUS OF THESE RECORDS

A brief and a record are the same object at two moments in its life, which is why the comparator, the estimand and the evidence level survive into the result. Every organisation, figure and estimate on these two pages has been invented, and none constitutes a claim about Coglytix performance.



What the evidence already shows

Let us examine fifteen documented cases across specific industries. Each case names an intervention, identifies its effect using a stated method, and reports a result and its limitations. Green indicates peer-reviewed sources, while navy indicates preprints, company-reported results and practitioner articles.

HEALTH	A cost proxy for need left the sickest patients unreferred [8]	Algorithm audit
CLIMATE	Human influence at least doubled the risk of a heatwave in the summer of 2003 that killed 70,000 [50], [51]	Counterfactual attribution
PUBLIC HEALTH	Twelve diarrhoea correlates narrowed to three drivers [8]	Causal Bayesian network
MATERNAL CARE	Delivery plans caused facility births; distance did not [52]	Causal ML on survey data
CLIMATE COMMS	Consensus messaging backfired only among a few US conservatives [53]	Bayesian network model
FUNDING	Resources placed where the causal effect is strongest [8]	Effect-ranked allocation
PAYMENTS	3% higher checkout conversion against a holdout [54]	Experimental holdout
ENERGY	7,000 homes cut about 0.10 kWh, or 11%, per price event [55]	Field experiment, heterogeneous effects
LENDING	Payday loans caused later defaults and overdrafts [56]	Fuzzy regression discontinuity
RETAIL	Two of five coupon types actually raised sales [57]	Causal ML, policy learning
MICROFINANCE	Group loans funded about 3.3 days faster on Kiva [58]	Treatment-effect estimation
MOBILITY	Incentives allocated by uplift, not propensity [59]	Uplift from experiment data
TRAVEL	66% of full-treatment uplift while treating 47% [60]	Uplift, online randomised trials
STREAMING	Localisation valued where withholding it is not an option [61]	DiD, double ML, synthetic control
MANUFACTURING	Fault analysis with fewer spurious causal links [62]	Causal graph, expert input

What the record shows and what it does not. Causal methods are already informing decisions in production through holdouts and randomised trials, quasi-experimental identification, uplift and heterogeneous-effect targeting, and causally structured root-cause analysis. What the public record does not show is the layer beneath them: shared definitions, decisions written as specifications, and signed records of what was decided and what followed. In each case above, that scaffolding was built once for a single decision.



Built around real business situations

FIGURE 20 · USE CASE MAP





From noise to impact


Coglytix is designed to move teams from confusion to clarity. The journey below tracks how a founder or operator experiences the shift.

FIGURE 21 · USER JOURNEY FROM NOISE TO IMPACT


CURRENT STATE



COGLYTIX EXPERIENCE



OUTCOME





Business value and product principles

FIGURE 22 · BUSINESS VALUE MODEL




FIGURE 23 · PRODUCT PRINCIPLES



The product is built to reduce decision fatigue and not create another complicated system.



09

Why now





Convergence and urgency

Three conditions have only recently converged, and a fourth adds urgency.

- 01** The causal toolkit has matured from a set of academic methods into deployable open-source libraries [21], [22].
- 02** Generative AI has made it practical to extract structured decision context from the unstructured sources where it has always been hidden.
- 03** The enterprise has largely built the two layers below, yet value lags investment: 99.1% of data and AI leaders call investment in data and AI a top priority, but only 54.0% report high or significant business value from it [63].
- 04** A fourth adds urgency: the EU AI Act is being phased in. Prohibitions and general-purpose AI rules already apply; high-risk duties apply from 2 December 2027, or from 2 August 2028 for AI in regulated products [44], [64].

The market has reached the limits of the analytics layer and discovered that analytics is not the top of the value chain. The layer above it, the decision layer, remains largely unbuilt. It is where data finally becomes impact.



The infrastructure imperative

The field's own account of what is required for wider adoption points to the need for permanent, reusable infrastructure rather than isolated projects [8]. Each requirement described below is therefore built into the decision layer.

- 01** Make better use of the data already collected. Existing datasets sit in separate systems at differing levels of granularity and can be joined only once their terms are agreed, which the ontology is designed to provide [8].
- 02** Collect more comprehensive data. Behaviour is shaped not only by what people believe and want, but also by institutions, rules and constraints. Data collected around fixed assumptions about what matters may therefore miss the true drivers entirely [8].
- 03** Build on open tooling. Proprietary platforms priced whole sectors out of causal work; open engines are cheaper, inspectable and improve faster through shared use [8], and bnlearn [65], DoWhy [21] and EconML [22] make this practical today.
- 04** Combine artificial and human intelligence. Experts constrain the model, identify confounders omitted from the data and interpret counterintuitive results. Ethical and policy judgements must remain within the same decision loop [8]
- 05** Improve how models are evaluated. Test each method against cases where the answer is already known, either through established findings or simulated data with a fixed causal structure [8]. Within the decision layer, the difference between predicted and realised effects provides a continuous test of the model's reliability.
- 06** Demonstrate value one decision at a time. Published examples remain limited [8], so every validated decision becomes a valuable asset. The signed Decision Memory captures and accumulates this evidence over time.
- 07** Build the literacy of the people responsible for making decisions. Leaders need a clear understanding of the concepts, not expertise in the mathematics. The two interfaces described in Section 06 are designed to make this possible [8]

Shared definitions, open engines, human review and measured outcomes are properties of a layer, not of a single analysis. That is the case for building the decision layer rather than commissioning another study.



10

The call to category
and the future vision





The call to category

Data infrastructure made data available. Analytics infrastructure made it understandable. Neither made decisions accountable. The organisations that succeed in the next decade will be those that treat each decision as the first-class, versioned and provable artefact it has always deserved to be. The infrastructure layer that makes this possible is the category defined in this paper.

The decision is the most valuable
thing a business produces.

It is time it had a home.



The future of business analytics & insights will not be limited to dashboards and score cards. **It will be decision centred.**

Businesses will need systems that remember what was decided, why it was decided, who acted, what happened and what should be learned. Coglytix is building towards a world where every important business decision has context, evidence, ownership, accountability and measurable learning attached to it.

FIGURE 24 · FUTURE STATE



Growth creates noise. Coglytix brings clarity.

It starts where analytics stops. It closes the gap between data and measurable impact. It helps growing businesses see clearly, decide confidently and grow intelligently.



11

Design partner programme





Help shape what we are building

Coglytix is inviting founders, operators, business leaders and decision-focused teams to help shape the product. The programme is for organisations that want to improve how they use data and AI to make better decisions.

DESIGN PARTNERS HELP SHAPE

- The most important decision use cases
- The data sources that matter most
- The decision briefs leaders need
- The metrics and drivers that should be tracked
- The workflow from data to insight to action
- The trust and explainability requirements
- The impact tracking experience

Become a Coglytix design partner and help shape the causally grounded decision intelligence infrastructure modern businesses need.

Complete our decision research survey @ coglytix.ai OR Contact the team: hello@coglytix.ai





AP

About this paper:
status, citation and
notices





About this paper

Status of the work

This paper sets out the architecture Coglytix is building and the case for it. Capabilities are described as designed rather than as delivered, and the worked specification in Section 06 is illustrative. Nothing here is a commitment to a delivery date or a specific implementation.

How to cite

Coglytix Ltd (2026). The Causal Decision Intelligence Infrastructure for Modern Businesses. Revised edition, version 1.0, September 2026. Available at coglytix.ai.

Sources and method

References are numbered in order of first citation and were checked against publisher, DOI or official records in September 2026. The applied cases in Section 08 are drawn from Sgaier, Huang and Charles [7]; each is attributed to the study it comes from.

Regulatory position

Statements about the EU AI Act and the NIST AI Risk Management Framework describe the position as at September 2026 and may change; the Act's high-risk timetable has already been amended once. Nothing in this paper is legal advice.

Forward-looking statements

Parts of this paper describe intentions and expectations rather than present facts. They are subject to change, and no reliance should be placed on them as a warranty, a representation or an offer.

Trademarks and contact

All third-party names and marks are the property of their owners, and no affiliation or endorsement is implied. Questions, corrections and design-partner enquiries: hello@coglytix.ai.



GL

Glossary of
working terms





Glossary of working terms

A

Ontology

A formal, shared model of the organisation's entities, relationships, and metrics, defined once, on which valid causal reasoning and reliable extraction depend.

Declarative Decisioning

Decisions expressed as declarative, version-controlled specifications; the code-level interface to the engine.

Decision Studio

A no-code, cloud-delivered visual interface to the same engine, in which decisions are composed through a guided UI that compiles to the same Decision Specification.

Decision Specification

The declared objective, assumptions, data, constraints, and policy of a decision.

Signed Decision Memory

The cryptographically signed, hash-chained, append-only store of decisions and their realised outcomes, tamper-evident and replayable for audit.

Plan, apply, track

The decision lifecycle: predict before committing, execute, and reconcile predicted against realised impact.



Glossary of working terms

B

Coglytix Core

The decision intelligence foundation layer beneath every Coglytix product, connecting data, context, causal reasoning, decisions, owners, actions and results in one accountable loop.

Decision brief

A structured decision object: issue, metrics, evidence, likely drivers, options, trade-offs, recommended action, owner, expected outcome, success measure and review date.

Decision memory

The record of what was decided, why, by whom, on what evidence, what action followed and what was learned. The signed Decision Memory is its tamper-evident form.

Trusted metric

A decision-ready signal: a metric that is reliable and meaningful, linked to what it affects, and clear about whether it calls for a decision.

Ladder of Causation

Three rungs of causal reasoning: association (seeing), intervention (doing) and counterfactual (imagining). Decisions live on the upper two rungs.

Next best action

The action that best advances the objective under the causal model, the constraints and the organisation's risk tolerance: the answer to the prescriptive question.



Glossary of working terms

C

Potential outcomes

The comparison between what happened to a unit and what would have happened had it not been exposed. Propensity-score matching reconstructs a control group where none exists.

Causal Bayesian network

A graph model that estimates relationships among all the variables in a dataset, which makes it a discovery method as well as a test of stated relationships.

Counterfactual attribution

Estimating how far an influence changed the probability of an outcome, by simulating the world with that influence and without it.

Proxy metric

A measurable stand-in for the outcome wanted. It is safe only where its causal link to that outcome holds for every group it is applied to.

Causal discovery

Deriving candidate causal structure from data rather than specifying it in advance, then constraining that structure with expert knowledge.

Precision allocation

Directing limited resources to the interventions with the strongest estimated causal effect, rather than the largest predicted volume.



RF

References





References [01]–[18]

The argument draws on established work in causal inference, decision theory, AI governance and infrastructure design, together with regulatory and industry sources. References are numbered in order of first citation; online sources were checked in September 2026.

- [1] Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A. and Fung, P. (2023). Survey of Hallucination in Natural Language Generation. *ACM Computing Surveys*, 55(12), Article 248. <https://doi.org/10.1145/3571730>
- [2] Lustig, I., Dietrich, B., Johnson, C. and Dziekan, C. (2010). The Analytics Journey. *Analytics Magazine*, November/December 2010, 11–18. INFORMS. <https://doi.org/10.1287/LYTX.2010.06.01>
- [3] Pearl, J. (2009). *Causality: Models, Reasoning, and Inference* (2nd ed.). Cambridge University Press. <https://doi.org/10.1017/CBO9780511803161>
- [4] Morris, K. (2021). *Infrastructure as Code: Dynamic Systems for the Cloud Age* (2nd ed.). O'Reilly Media. Available at <https://www.oreilly.com/library/view/infrastructure-as-code/9781098114664/>
- [5] Pearl, J. and Mackenzie, D. (2018). *The Book of Why: The New Science of Cause and Effect*. Basic Books.
- [6] Bareinboim, E., Correa, J. D., Ibeling, D. and Icard, T. (2022). On Pearl's Hierarchy and the Foundations of Causal Inference. In H. Geffner, R. Dechter and J. Y. Halpern (eds), *Probabilistic and Causal Inference: The Works of Judea Pearl*, 507–556. ACM Books. <https://doi.org/10.1145/3501714.3501743>
- [7] Bender, E. M., Gebru, T., McMillan-Major, A. and Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FACt '21)*, 610–623. <https://doi.org/10.1145/3442188.3445922>
- [8] Sgaier, S. K., Huang, V. and Charles, G. (2020). The Case for Causal AI. *Stanford Social Innovation Review*, 18(3). <https://doi.org/10.48558/KT81-SN73>
- [9] Obermeyer, Z., Powers, B., Vogeli, C. and Mullainathan, S. (2019). Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations. *Science*, 366(6464), 447–453. <https://doi.org/10.1126/science.aax2342>
- [10] Manson, J. E., Cook, N. R., Lee, I.-M. et al. (2019). Vitamin D Supplements and Prevention of Cancer and Cardiovascular Disease. *New England Journal of Medicine*, 380(1), 33–44. <https://doi.org/10.1056/NEJMoa1809944>
- [11] Pittas, A. G., Dawson-Hughes, B., Sheehan, P. et al. (2019). Vitamin D Supplementation and Prevention of Type 2 Diabetes. *New England Journal of Medicine*, 381(6), 520–530. <https://doi.org/10.1056/NEJMoa1900906>
- [12] The Royal Swedish Academy of Sciences (2019). The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 2019: Abhijit Banerjee, Esther Duflo and Michael Kremer, “for their experimental approach to alleviating global poverty”. <https://www.nobelprize.org/prizes/economic-sciences/2019/summary/>
- [13] Pfeffer, J. and Sutton, R. I. (2000). *The Knowing–Doing Gap: How Smart Companies Turn Knowledge into Action*. Harvard Business School Press. ISBN 978-1-57851-124-2.
- [14] Ackoff, R. L. (1989). From Data to Wisdom: Presidential Address to ISGSR, June 1988. *Journal of Applied Systems Analysis*, 16, 3–9
- [15] Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W. and Robins, J. (2018). Double/Debiased Machine Learning for Treatment and Structural Parameters. *The Econometrics Journal*, 21(1), C1–C68. <https://doi.org/10.1111/ectj.12097>
- [16] Kohavi, R., Tang, D. and Xu, Y. (2020). *Trustworthy Online Controlled Experiments: A Practical Guide to A/B Testing*. Cambridge University Press. <https://doi.org/10.1017/9781108653985>
- [17] Angrist, J. D. and Pischke, J.-S. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.
- [18] Abadie, A., Diamond, A. and Hainmueller, J. (2010). Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program. *Journal of the American Statistical Association*, 105(490), 493–505. <https://doi.org/10.1198/jasa.2009.ap08746>



References [19]–[36]

- [19] Wager, S. and Athey, S. (2018). Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. *Journal of the American Statistical Association*, 113(523), 1228–1242. <https://doi.org/10.1080/01621459.2017.1319839>
- [20] Imbens, G. W. and Rubin, D. B. (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction* (ch. 22, Sensitivity Analysis and Bounds). Cambridge University Press. <https://doi.org/10.1017/CBO9781139025751>
- [21] Sharma, A. and Kiciman, E. (2020). DoWhy: An End-to-End Library for Causal Inference. arXiv:2011.04216. <https://arxiv.org/abs/2011.04216>
- [22] Battocchi, K., Dillon, E., Hei, M., Lewis, G., Oka, P., Oprescu, M. and Syrgkanis, V. (2019). EconML: A Python Package for ML-Based Heterogeneous Treatment Effects Estimation [software]. PyWhy. <https://github.com/py-why/EconML>
- [23] Morgan, S. L. and Winship, C. (2015). *Counterfactuals and Causal Inference: Methods and Principles for Social Research* (2nd ed.). Cambridge University Press. <https://doi.org/10.1017/CBO9781107587991>
- [24] Howard, R. A. (1966). Information Value Theory. *IEEE Transactions on Systems Science and Cybernetics*, 2(1), 22–26. <https://doi.org/10.1109/TSSC.1966.300074>
- [25] Rosenbaum, P. R. and Rubin, D. B. (1983). The Central Role of the Propensity Score in Observational Studies for Causal Effects. *Biometrika*, 70(1), 41–55. <https://doi.org/10.1093/biomet/70.1.41>
- [26] Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann.
- [27] Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J. and Amodei, D. (2020). Scaling Laws for Neural Language Models. arXiv:2001.08361. <https://arxiv.org/abs/2001.08361>
- [28] The Payment Services Regulations 2017, SI 2017/752, regulation 71 (limits on the use of payment instruments), regulation 82 (refusal of payment orders) and regulation 100 (authentication). <https://www.legislation.gov.uk/uksi/2017/752/regulation/100>
- [29] Financial Conduct Authority (2021). *Technical Standards 2021: Strong Customer Authentication and Common and Secure Methods of Communication*, articles 2, 4, 5 and 18 to 20 (transaction risk analysis). FCA Handbook. <https://handbook.fca.org.uk/technical-standards/s140c1223>
- [30] The Payment Services (Amendment) Regulations 2024, SI 2024/1013, in force 30 October 2024, inserting regulation 86(2A) to (2D) of SI 2017/752 (delay of a payment order on reasonable grounds of suspected fraud). <https://www.legislation.gov.uk/uksi/2024/1013/body/made>
- [31] Financial Conduct Authority (2024). FG24/6: Guidance for firms that enables a risk-based approach to payments, 22 November 2024. <https://www.fca.org.uk/publications/finalised-guidance/fg24-6-guidance-firms-enables-risk-based-approach-payments>
- [32] Payment Systems Regulator (2025). APP Scams Reimbursement Requirement: Consolidated Policy Statement, PS25/5, May 2025, giving effect to Specific Direction 20 for Faster Payments, in force 7 October 2024. <https://www.psr.org.uk/media/rhelv4op/ps25-5-app-scams-reimbursement-consolidated-policy-statement-may-2025.pdf>
- [33] Ofgem (2025). Standard Conditions of Electricity Supply Licence, Standard Licence Condition 27.8 (ability to pay), consolidated 1 August 2025. <https://www.ofgem.gov.uk/sites/default/files/2025-08/Electricity-Supply-Standard-Consolidated-Licence-Conditions.pdf>
- [34] The Offshore Installations (Offshore Safety Directive) (Safety Case etc.) Regulations 2015, SI 2015/398, regulation 7 (corporate major accident prevention policy), regulation 17 (safety case for a production installation) and regulations 23 to 24 (review and revision). <https://www.legislation.gov.uk/uksi/2015/398/regulation/17/made>
- [35] The Pressure Systems Safety Regulations 2000, SI 2000/128, regulation 8 (written scheme of examination) and regulation 12 (maintenance in good repair), read with the definition of pipework in regulation 2, which includes compressors. <https://www.legislation.gov.uk/uksi/2000/128/regulation/12/made>
- [36] The Offshore Installations and Pipeline Works (Management and Administration) Regulations 1995, SI 1995/738, regulation 10 (permits to work). <https://www.legislation.gov.uk/uksi/1995/738/made>



References [37]–[54]

- [37] Financial Conduct Authority. CONC 7.3.4R, Treatment of customers in or approaching arrears or in default. FCA Handbook, Consumer Credit sourcebook. <https://www.handbook.fca.org.uk/handbook/CONC/7/3.html>
- [38] Ofcom (2026). General Conditions of Entitlement, General Condition C1 (contract requirements, including end-of-contract notification) and C7 (switching), effective 8 April 2026. <https://www.ofcom.org.uk/phones-and-broadband/accessibility/general-conditions-of-entitlement>
- [39] Ofcom (2026). Investigation into Virgin Media’s compliance with contract termination and complaints handling rules, case CW/01275/07/23, Confirmation Decision 8 July 2026. <https://www.ofcom.org.uk/siteassets/resources/documents/about-ofcom/bulletins/enforcement-bulletin/virgin-media/virgin-media-difficulty-cancelling-non-confidential-confirmation-decision-96C.pdf?v=423911>
- [40] Financial Conduct Authority. ICOBS 6B.2.1R, Setting renewal prices (ICOBS 6B, Home insurance and motor insurance pricing). FCA Handbook. <https://www.handbook.fca.org.uk/handbook/ICOBS/6B/2.html>
- [41] Crosby, S. A. and Wallach, D. S. (2009). Efficient Data Structures for Tamper-Evident Logging. Proceedings of the 18th USENIX Security Symposium. USENIX Association. https://www.usenix.org/legacy/event/sec09/tech/full_papers/crosby.pdf
- [42] National Institute of Standards and Technology (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1. U.S. Department of Commerce. <https://doi.org/10.6028/NIST.AI.100-1>
- [43] Pratt, L. (2019). Link: How Decision Intelligence Connects Data, Actions, and Outcomes for a Better World. Emerald Publishing. <https://doi.org/10.1108/9781787696532>
- [44] European Parliament and Council of the European Union (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, OJ L, 2024/1689, 12 July 2024. <http://data.europa.eu/eli/reg/2024/1689/oj>
- [45] Gutierrez, P. and Gérardy, J.-Y. (2017). Causal Inference and Uplift Modelling: A Review of the Literature. Proceedings of the 3rd International Conference on Predictive Applications and APIs, PMLR 67, 1–13. <https://proceedings.mlr.press/v67/gutierrez17a.html>
- [46] Kusner, M. J., Loftus, J., Russell, C. and Silva, R. (2017). Counterfactual Fairness. Advances in Neural Information Processing Systems 30 (NeurIPS 2017). <https://arxiv.org/abs/1703.06856>
- [47] Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A. and Bengio, Y. (2021). Toward Causal Representation Learning. Proceedings of the IEEE, 109(5), 612–634. <https://doi.org/10.1109/JPROC.2021.3058954>
- [48] Nabi, R. and Shpitser, I. (2018). Fair Inference on Outcomes. Proceedings of the AAAI Conference on Artificial Intelligence, 32(1), 1931–1940. <https://doi.org/10.1609/aaai.v32i1.11553>
- [49] Chiappa, S. (2019). Path-Specific Counterfactual Fairness. Proceedings of the AAAI Conference on Artificial Intelligence, 33(01), 7801–7808. <https://doi.org/10.1609/aaai.v33i01.33017801>
- [50] Stott, P. A., Stone, D. A. and Allen, M. R. (2004). Human Contribution to the European Heatwave of 2003. Nature, 432(7017), 610–614. <https://doi.org/10.1038/nature03089>
- [51] Robine, J.-M., Cheung, S. L. K., Le Roy, S., Van Oyen, H., Griffiths, C., Michel, J.-P. and Herrmann, F. R. (2008). Death Toll Exceeded 70,000 in Europe during the Summer of 2003. Comptes Rendus Biologies, 331(2), 171–178. <https://doi.org/10.1016/j.crv.2007.12.001>
- [52] Huang, V. S., Morris, K., Jain, M. et al. (2020). Closing the Gap on Institutional Delivery in Northern India: A Case Study of How Integrated Machine Learning Approaches Can Enable Precision Public Health. BMJ Global Health, 5(10), e002340. <https://doi.org/10.1136/bmjgh-2020-002340>
- [53] Cook, J. and Lewandowsky, S. (2016). Rational Irrationality: Modeling Climate Change Belief Polarization Using Bayesian Networks. Topics in Cognitive Science, 8(1), 160–179. <https://doi.org/10.1111/tops.12186>
- [54] Fial de Sá, F. (2024). Unified Checkout: Streamlining Uber’s Payment Ecosystem. Uber Engineering Blog, 14 November 2024. <https://www.uber.com/blog/unified-checkout/>



References [55]–[65]

- [55] Zhou, D. P., Balandat, M. and Tomlin, C. J. (2017). Estimating Heterogeneous Treatment Effects in Residential Demand Response. arXiv:1710.03190. <https://arxiv.org/abs/1710.03190>
- [56] Gathergood, J., Guttman-Kenney, B. and Hunt, S. (2019). How Do Payday Loans Affect Borrowers? Evidence from the U.K. Market. *The Review of Financial Studies*, 32(2), 496–523. <https://doi.org/10.1093/rfs/hhy090>
- [57] Langen, H. and Huber, M. (2023). How Causal Machine Learning Can Leverage Marketing Strategies: Assessing and Improving the Performance of a Coupon Campaign. *PLOS ONE*, 18(1), e0278937. <https://doi.org/10.1371/journal.pone.0278937>
- [58] Pham, T. T. and Shen, Y. (2017). A Deep Causal Inference Approach to Measuring the Effects of Forming Group Loans in Online Non-profit Microfinance Platform. arXiv:1706.02795. <https://arxiv.org/abs/1706.02795>
- [59] Tulasyan, R. and Aggarwal, B. (2026). Beyond Prediction: Solving the Multiple Knapsack Problem at Scale: How Uber Optimizes Incentives. Uber Engineering Blog, 14 May 2026. <https://www.uber.com/blog/solving-multiple-knapsack/>
- [60] Goldenberg, D., Albert, J., Bernardi, L. and Estevez, P. (2020). Free Lunch! Retrospective Uplift Modeling for Dynamic Promotions Recommendation within ROI Constraints. *Proceedings of the 14th ACM Conference on Recommender Systems (RecSys '20)*, 486–491. <https://doi.org/10.1145/3383313.3412215>
- [61] Netflix Technology Blog (2022). A Survey of Causal Inference Applications at Netflix, 21 May 2022. <https://netflixtechblog.com/a-survey-of-causal-inference-applications-at-netflix-b62d25175e6f>
- [62] Wehner, C., Kertel, M. and Wewerka, J. (2023). Interactive and Intelligent Root Cause Analysis in Manufacturing with Causal Bayesian Networks and Knowledge Graphs. 2023 IEEE 97th Vehicular Technology Conference (VTC2023–Spring). <https://doi.org/10.1109/VTC2023-Spring57618.2023.10199563>
- [63] Bean, R. (2026). 2026 AI & Data Leadership Executive Benchmark Survey (Exhibits D and W), released December 2025. Data & AI Leadership Exchange. <https://static1.squarespace.com/static/62adf3ca029a6808a6c5be30/t/6942c3cb535da44088c2dbff/1765983179572/2026+AI+%26+Data+Leadership+Executive+Benchmark+Survey+Final.pdf>
- [64] European Parliament and Council of the European Union (2026). Regulation (EU) 2026/1744 of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI). *Official Journal of the European Union*, OJ L, 2026/1744, 24 July 2026. <https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng>
- [65] Scutari, M. (2010). Learning Bayesian Networks with the bnlearn R Package. *Journal of Statistical Software*, 35(3), 1–22. <https://doi.org/10.18637/jss.v035.i03>



THE CAUSAL DECISION INTELLIGENCE INFRASTRUCTURE FOR MODERN BUSINESSES

Coglytix is the decision layer **for profitable growth**

Turn data noise into decisions that drive the desired business outcomes.

BECOME A DESIGN PARTNER

Message us if you would like to **join our design partner programme**

hello@coglytix.ai · coglytix.ai