A FREE EXECUTIVE GUIDE

AI Integration, Mathematical Modelling, and Governance

The Nessiva Approach to Better Business Decisions and Performance

Maksim Sokolov
Sanjar Alimov

Co-founders of Nessiva



nessiva.com

Updated edition · October 2026

About this guide

Free digital edition · October 2026

Copyright © 2026 Maksim Sokolov and Sanjar Alimov

An executive guide from Nessiva

Available free of charge at nessiva.com

Written by the co-founders of Nessiva, this guide explains how AI integration, mathematical modelling, governance, and cybersecurity can support business performance. It also introduces the services our team provides.

The discussion presents management principles without disclosing proprietary formulas, algorithms, calibration methods, source code, or confidential client information.

Implementation experience is identified separately from work in development and illustrative applications. Numerical examples are hypothetical unless stated otherwise. No client endorsement is implied.

Product descriptions reflect the documentation reviewed in September 2026. Specific capabilities and availability may change. References to third-party products and frameworks do not imply affiliation or endorsement.

This October 2026 update introduces the *Virgilio* and *Dante* model names, the three kernel designs, and the three levels of engagement. The existing dated product references retain their original review dates.

Contents

For the business leader....iv

1 The language of this book....1

2 An established discipline with new economics....8

3 What the latest AI capabilities make possible....10

4 Define the benefit before choosing the tool....13

5 What mathematical modelling adds....15

6 The mathematical skeleton....18

7 Progressive model development....23

8 A business model that lives over time....27

9 Traces and reliable calculations....30

10 Where AI improves everyday work....33

11 Make AI part of a controlled workflow....37

12 Connect the business through ChatGPT....40

13 Agents and kernels working together....44

14 Governance makes the system accountable....48

15 Cybersecurity for connected AI....53

16 The Risk and AI Officer function....58

17 Measure value in the complete process....60

18 Begin with a practical sequence....63

19 Work with Nessiva....66

About the authors....70

Sources and notes....72

For the business leader

A business can be busy and profitable while losing value in ways that are difficult to see. People search for information that should be available. A delay in one department puts another under pressure. A decision that improves sales today can make tomorrow's commitments harder to meet.

This guide explains how artificial intelligence, mathematical modelling, and governance can help address these problems together. Mathematics makes selected business relationships explicit. AI accelerates the preparation, execution, and explanation of the work. Governance establishes authority and accountability, while cybersecurity protects the information and systems involved.

The disciplines are established, and our approach draws on professional experience developed over many years. Recent AI capabilities make selected analytical and operating tools more accessible to smaller businesses. The opportunity is to improve useful work now while developing a more informed way to manage the business over time.

We are the co-founders of Nessiva, and we offer this guide free of charge to help business leaders assess these possibilities. It also introduces our services in mathematical modelling, AI integration, risk management, governance, and cybersecurity. Our aim is to give you a useful basis for evaluating opportunities and deciding whether to work with our team.

We begin with definitions, then develop the argument from business outcomes to the mathematical skeleton, connected AI, and evidence over time. Governance, cybersecurity, and practical implementation complete the discussion. The closing chapter explains how a business can begin working with Nessiva.

You do not need to understand the formulas. You need to understand the business question, the evidence supporting an answer, and the responsibilities surrounding its use.

Maksim Sokolov and Sanjar Alimov

CHAPTER 01

The language of this book

Clear definitions prevent expensive misunderstandings. The terms below have the following meanings throughout this book. They are working definitions for a business audience. Type A, Type A-B, and AI kernel describe our approach; they are not universal industry classifications.

Business results and measurement

Business outcome. A result the organization values, such as healthy cash flow, dependable delivery, customer satisfaction, technical quality, or the ability to keep operating through disruption. An activity is something people do; an outcome is what the business gains or loses as a result.

Performance. How well an activity or organization meets its agreed requirements and contributes to valued outcomes. Performance includes quality and reliability as well as speed. Finishing quickly does not establish good performance if the work has to be repeated.

Capacity. The amount and mix of work an operation can sustain over a stated period while meeting its agreed quality and service requirements under specified working conditions. Capacity depends on available skills, equipment, support, and the nature of the work.

Workload. The work requiring an operation's attention during a stated period. **Resource consumption** is what the work actually uses, such as staff hours or machine time. **Operational strain** is the pressure associated with meeting that workload, visible through evidence such as persistent overtime, growing queues, or postponed checks. These concepts describe different things.

Constraint. A limit that a decision must respect, such as a delivery obligation, a specialist's availability, or a quality requirement. A **bottleneck** is the part of a process that restricts its overall flow under the conditions being examined.

Evidence. Information with an identifiable basis that can support or challenge a conclusion. Records, observations, and documented expert assessments can all be evidence, but they have different strengths and limitations. An **assumption** is something provisionally accepted for an analysis and kept visible for review.

Measure. A defined way to describe an aspect of the business using observations or numbers. A **baseline** is the recorded starting position against which a proposed improvement is compared. **Validation** is the examination of whether a measure, model, or tool is sufficiently sound and useful for its intended purpose.

Mathematical modelling

Quantification. Expressing relevant aspects of the business in defined, comparable measures. It requires deciding what a number represents, how it is obtained, and what conclusions it can support. Giving an opinion a numerical label does not automatically make it objective.

Mathematical model. An explicit representation of selected relationships, assumptions, and constraints, expressed through rules that can be calculated. In business, a model helps organize evidence and compare the consequences of choices. It is a purposeful simplification of the business.

Business stream. A recognizable flow of work and value, such as a product line, project category, or service offering. A **capability** is an enduring ability needed to support that work, such as selling, design, production, delivery, or customer support. Several streams may depend on the same capability.

Utility vector. The set of business outcome measures that leadership has chosen to keep visible together. Here, utility means value to the business. The word vector means that several distinct outcomes are considered together instead of being hidden inside one number.

Mathematical skeleton. Our description of the durable structure connecting business outcomes, operating relationships, measures, assumptions, and constraints. It gives separate tools a common analytical purpose and preserves meaning across repeated use. It can change through an explicit review process.

Virgilio. The name of Nessiva's structured business-health model, built around the Utility Vector Framework. It brings organizational evidence into a common structure and connects

performance, capacity and risk to valued outcomes. Its staged development includes the Type A foundation and deeper Type A-B work.

Dante. The collective name for Nessiva's data-analytics models. These use structured evidence from *Virgilio* and other relevant, validated data to investigate correlations, regressions, trends and delayed effects. Suitable forecasting and scenario methods may extend the analysis when the evidence supports them.

Type A. The foundation of the Nessiva framework: a company-specific view connecting valued outcomes, business streams, and organizational capabilities. It establishes a common structure for examining performance, capacity, and risk.

Type A-B. A deeper investigation within selected parts of the Type A structure. It connects operating processes and evidence to the broader management view. Our website writes this as Type A.B; Type A-B means the same stage in this book.

Parameter. A setting that affects how a model behaves, such as an approved priority or a relationship used in calculation. **Calibration** means selecting and adjusting such settings using appropriate evidence and judgment. These settings require controlled review because changing them can change the meaning of the results.

Scenario analysis. Comparing what a model implies under different stated conditions. A **stress test** examines demanding or adverse conditions. A **simulation** calculates outcomes for many constructed cases; a Monte Carlo simulation varies inputs using specified random patterns. The usefulness of any of these methods depends on the assumptions behind the cases.

Time-series analysis. Examining comparable observations collected over time. **Regression analysis** is a statistical method for estimating relationships among measured factors. **Forecasting** uses evidence and assumptions to estimate future outcomes. A statistical relationship alone does not establish that one factor causes another.

Snapshot. A record of a position at a stated time. A **reporting period** is the interval described by a report. A **time stamp** identifies when an event or record occurred. These distinctions prevent activity during a period from being confused with a position at its end.

Trend. A sustained direction of movement in a series of observations. **Seasonality** is a recurring pattern associated with a

calendar or operating cycle. A **lag** is a delay between related events or effects. A **trajectory** describes the path of a measure, or a set of measures, through time.

Trace. A preserved record linking relevant inputs, methods, outputs, decisions, and later events. It supports examination of what happened. A **decision log** records a decision, its basis, owner, and intended result. Neither term refers to a language model's hidden internal reasoning.

Artificial intelligence in business

Artificial intelligence or AI. Computer systems that use learned patterns and other methods to perform tasks such as interpreting information, identifying patterns, generating content, or supporting decisions. The term covers different technologies with different capabilities.

Generative AI. AI that produces content, including text, images, and code, in response to information and instructions. Much of the AI discussed in this book is generative AI used alongside conventional software and calculation tools.

AI integration. Incorporating AI into an actual business process with defined inputs, responsibilities, access permissions, review requirements, and measures of success. Buying access to an AI service is one possible starting step; integration concerns how the work is performed.

Workflow. The sequence of tasks, decisions, and handoffs through which work is completed. **Automation** allows software to perform specified parts of that sequence. Automation may use AI, conventional rules, or both.

AI assistant. A system that helps a person perform a task, often through conversation. An **AI agent** is a system configured to work towards an assigned objective through multiple steps and, where permitted, use tools or take actions. The agent's authority must be defined separately from its technical capability.

AI kernel. In our terminology, a controlled package of instructions, definitions, reference knowledge, and operating rules that guides AI within a particular business or analytical context. It is not a guarantee of correct behaviour or a substitute for software security controls.

Kernel types. Nessiva distinguishes autonomous, deterministic and judgment-embedded kernels. Autonomous kernels coordinate permitted tasks within a defined mandate. Deterministic kernels execute explicit calculation rules. Judgment-embedded kernels place AI-assisted interpretation at specified points within a controlled process. These are Nessiva working categories, and a system may combine them.

Reusable skill. A documented, repeatable procedure that an AI system can use for a defined task, such as preparing a report from approved records. A kernel supplies common context and rules; a skill describes how a particular task is carried out.

Company configuration. Preparing an AI environment with approved business documents, terminology, instructions, and access. This does not necessarily change or retrain the underlying AI model. In this book, an AI model is the general technology; a business model is the company-specific analytical representation.

ChatGPT and an AI model. ChatGPT is a product through which people use AI capabilities. An underlying model, such as GPT-6 Astra, supplies reasoning and generation capabilities. The surrounding tools, permissions, and operating environment determine what the complete system can access and do.

Business workspace. An organization's managed environment for users and shared resources. A **project** groups related work and context. A **custom GPT** is a reusable assistant configuration for a defined purpose. These concepts describe different ways to organize work; none automatically provides a complete business record.

Plugin and connection. A plugin extends an AI environment with supported capabilities, such as reusable skills or access to an external service. A connection, sometimes called an app or connector, provides authorized access to specified information or actions. An **API** is a defined software interface through which systems can interact.

Scheduled task and event trigger. A scheduled task repeats work at specified times. An event trigger starts work when a supported event occurs. Both require an executable procedure and accessible inputs. **Durable storage** preserves required records beyond one conversation or execution.

Responsibility and protection

Risk. The possibility that uncertain events or decisions will adversely affect valued business outcomes. The loss may involve money, time, quality, trust, security, or future capability. Risk management identifies, assesses, and addresses those possibilities.

Governance. The allocation of authority, accountability, and oversight. It establishes who may decide, who must review, and how concerns are resolved. A **policy** states an organizational rule or expectation. A **control** is a practical measure that helps put that rule into effect.

Cybersecurity. Protecting digital systems and information against unauthorized access, disruption, alteration, and misuse. **AI security** applies this concern to AI systems and their connections, including the possibility that untrusted content could manipulate an AI workflow.

Prompt injection. An attempt to make an AI system follow unauthorized instructions placed in information it receives, including material it is supposed to review. A document's content should not acquire the authority to change the workflow's rules.

Model risk. The possibility of harm from a model that is unsuitable, incorrectly implemented, poorly maintained, misunderstood, or used beyond its intended scope. **Human oversight** means that named people have the information, competence, time, and authority to review, challenge, and intervene.

Least privilege. Giving a person or system only the access required for its authorized task. **Information integrity** means preserving records and calculations against improper alteration. **Operational resilience** is the ability to sustain or restore important work during disruption.

Assurance. A structured examination of whether evidence supports confidence in a process or control, within a stated scope. A review is different from formal certification or attestation. **Residual risk** is the risk remaining after the selected controls have been applied.

Pilot. A limited trial designed to answer whether a proposed workflow is useful and sufficiently controlled. **Scaling** means extending a successful approach to more work, people, or business areas after its performance and controls have been assessed.

Risk and AI Officer. A function that coordinates quantitative analysis, selected AI applications, and the rules governing their use. A **fractional** appointment provides this function for an agreed portion of professional time, with a defined mandate and reporting cycle.

CHAPTER 02

An established discipline with new economics

Businesses have used mathematical analysis, statistics, forecasting, process improvement, and risk management for many decades. The questions are familiar: how much stock to carry, where to assign scarce resources, which commitments are sustainable, and what to do when conditions change. The arrival of generative AI did not create these questions or the disciplines used to answer them.

Our approach combines this established foundation with years of professional experience in quantitative reasoning, technology assurance, and governance.

What has changed is the practical effort required to turn expertise into working tools. A useful analytical model requires more than an idea. Someone must organize the data, implement the calculations, test the software, prepare explanations, maintain documentation, and repeat the analysis as the business changes. For an ambitious model, that supporting work can be substantial.

Why scale used to matter so much

Large banks and corporations have generally been better placed to sustain teams of analysts, quantitative specialists, programmers, and risk professionals. Their size can justify a continuing investment in tools and specialist roles. Smaller businesses have also used sophisticated analysis, but often with narrower scope or less support for ongoing development.

The historical barrier was therefore often organizational and economic. A medium-sized manufacturer might have a real need to understand interactions among design capacity, production commitments, cash, and customer obligations, yet be unable to justify a permanent analytical department around that need.

Modern AI can reduce parts of this burden. Under expert direction, it can assist with software development, document

preparation, evidence organization, testing, and repeated analytical tasks. It can help an expert prepare alternative implementations, examine edge cases, and communicate findings in language that managers can use.

Small and medium-sized businesses can therefore access selected advanced modelling capabilities without maintaining a large internal quantitative and programming department. A compact expert team can support a carefully scoped system. The cost depends on the data, complexity, integrations, and consequences of error. Competent people remain essential.

Two distinct sources of value

The first source of value is stronger analytical capability. AI helps qualified specialists build and operate models that might otherwise take more supporting effort. A business gains a more systematic way to examine dependencies and compare decisions.

The second source is better work within individual processes. AI can help prepare a quotation, compare document revisions, retrieve an approved procedure, or assemble a management report. Here the immediate benefit may be less repetitive work, faster preparation, or earlier identification of inconsistencies.

These sources of value reinforce each other. Better operational information can improve the management model. The model can indicate which operational improvement deserves attention first. The business then has a reason for choosing an AI application beyond the fact that the application is technically impressive.

What research supports

A study by Erik Brynjolfsson, Danielle Li, and Lindsey Raymond examined AI assistance among 5,172 customer-support agents. The revised working paper reported an average increase of 15 percent in issues resolved per hour, with substantial variation among workers. Less experienced workers benefited more; highly experienced workers saw smaller speed gains and some declines in quality. [1]

The study demonstrates possible gains in a specific setting, not a forecast for your business. Evaluate each proposed workflow on its own evidence, examining quality alongside speed.

CHAPTER 03

What the latest AI capabilities make possible

An executive who last tried AI as a conversational question-and-answer tool may underestimate what a properly equipped system can now do. The opportunity has expanded from receiving a useful answer to delegating a defined body of work: gathering approved information, applying a procedure, using calculation tools, producing a deliverable, and returning it for review.

This distinction matters economically. Much professional work consists of several connected steps. A good first draft helps, but a prepared comparison supported by source documents may remove a larger burden. The relevant question is how much of the complete process can be performed to an acceptable standard, with appropriate human involvement.

Astra and a broader choice of models

As of September 2026, OpenAI describes GPT-6 Astra as its most capable model for complex work across code, applications, and research, with advanced reasoning, computer use, and stronger judgment. Its documentation also distinguishes GPT-5.6 Sol for demanding work, Terra for everyday work, and Luna for clear, repeatable tasks. Access to individual models depends on the plan, product, and rollout. [2]

For a business leader, this creates a choice about how to allocate AI capability. A difficult investigation across several sources may justify a more capable model. A well-defined extraction task may be served by a faster, less expensive option after testing. The most capable available model need not perform every step.

Our practical interpretation is that model choice should become part of process design. Identify which tasks require judgment, which require exact calculation, and which mainly require repetition. Then select and test the resources appropriate to each. Spending more on reasoning can be justified when it improves a

consequential result; it can be unnecessary when the work is already tightly specified.

From conversation to finished work

Current ChatGPT documentation describes working with files, analysing data, producing documents and spreadsheets, using connected tools, and carrying out longer tasks. Projects can organize related chats, files, and instructions. Work mode supports substantial tasks that use tools and return outputs for review. These capabilities depend on the enabled environment and available permissions. [3]

Consider an illustrative procurement comparison. The useful output is not merely a paragraph naming a preferred supplier. It is a comparison with the relevant quotations, stated assumptions, missing information, and the business criteria applied. A manager should be able to inspect why an option appears attractive and what remains unresolved.

AI can assist in assembling that package. A mathematical model can provide the approved comparison logic. A purchasing manager can judge whether the evidence reflects commercial reality. Together, these contributions can shorten preparation while making the decision more explicit.

More capability changes the bottleneck

When software becomes easier to produce, choosing the right software becomes more important. When reports become easier to generate, deciding which reports deserve attention becomes more important. When agents can perform more steps, defining the limits and completion conditions of those steps becomes more important.

A business may discover that the limiting factor is no longer the effort of writing a report. It is agreement on what a completed job means, which record is authoritative, or what management should do when delivery reliability deteriorates. These are business design questions. Greater AI capability makes them harder to postpone.

This is why an impressive demonstration is only an opening. A system may handle the selected demonstration documents well and still struggle with incomplete scans, conflicting revisions, unusual contract structures, or missing operating data. The

business needs examples that represent its actual work, including difficult cases.

Improvement requires an acceptance decision

A newer model may perform better overall while changing a familiar workflow in ways that need attention. It may produce different explanations, use tools differently, or make different choices about when to ask a question. Existing procedures should therefore be checked before the new model takes over consequential work.

Keep a small collection of representative cases with known expectations. Include ordinary work, ambiguous inputs, missing evidence, and cases that should be referred to a person. Evaluate whether the new configuration preserves the required behaviour and improves the outcomes that justified the change.

The business should retain its approved definitions and computational method while testing the change in AI capability. This makes the comparison meaningful. If the model, source data, measurement rules, and review process all change together, management may struggle to explain why the results differ.

Build for continuing progress

Specific model names will change. A useful investment should preserve the company's knowledge, measures, records, and decision responsibilities as those names change. This does not mean every technical component is interchangeable without effort. It means that the business should own a clear description of what its system is intended to do.

Nessiva's emphasis is on that durable business structure. We can take advantage of stronger AI as it becomes suitable while maintaining continuity in the meaning of the work. The next chapters explain why this structure needs explicit mathematical relationships as well as clear operating instructions.

CHAPTER 04

Define the benefit before choosing the tool

An AI initiative becomes easier to evaluate when management can complete a plain sentence: we want to improve this part of the business because it is affecting this outcome.

For example, a manufacturer may want to shorten the preparation of production documents because incomplete information is delaying release. A retailer may want a clearer view of replenishment decisions because stock is accumulating in slow-moving items while customers cannot obtain popular ones. A service firm may want faster access to approved guidance because experienced staff are repeatedly answering the same internal questions.

Each statement connects work to a consequence. That connection should guide the choice of technology, the design of the workflow, and the measure of success.

Time creates value when it changes something

Saving staff time can be valuable in several ways. A team may reduce overtime, clear a backlog, give customers a faster answer, or assign skilled people to work that previously had to wait. These are different benefits and should be described separately.

If a person saves time but no useful work changes and no cost is avoided, the immediate financial result may be limited. The saving can still improve working conditions or provide capacity for future demand. Management should name that benefit accurately rather than treating every saved hour as cash in the bank.

The location of the saving matters as much as its size. Releasing time from a scarce reviewer may help the entire process move. Releasing time in an area that is already waiting for another department may have little effect on final delivery. This is one reason to connect local AI applications to a broader view of the business.

Quality can improve through earlier attention

Many information problems become expensive because they are discovered late. A missing reference might take minutes to resolve before production begins and much longer after work has proceeded on the wrong basis. An AI workflow that directs attention to the issue early can therefore have value beyond the time spent checking the document.

The mechanism must be explicit. What discrepancy can the system identify? What evidence will it show? Who will decide whether the discrepancy matters? What happens before the work moves forward? The benefit comes from this complete sequence.

Consistency is another potential gain. A controlled procedure can ask the same required questions across repeated cases, preserve the evidence reviewed, and make omissions easier to notice. AI can still miss information or produce unsupported conclusions. Consistency must be assessed in the workflow rather than assumed from the use of software.

Better decisions preserve several outcomes

Leaders rarely pursue one outcome in isolation. A rush order may generate attractive revenue while creating pressure on existing commitments. A lower-cost supplier may change quality or delivery risk. A decision to hold less inventory may release cash while reducing resilience to disruption.

The business needs a way to keep these consequences visible together. It may deliberately accept a sacrifice, but the sacrifice should be understood. A decision is stronger when management can explain what it expects to gain, what may deteriorate, and what conditions would justify changing course.

Knowledge becomes easier to use

Important knowledge is often scattered across people, documents, and past decisions. A controlled AI assistant can make approved information accessible to staff and identify unresolved questions. This is particularly useful during growth or coordination across locations. Source documents still need owners and review dates: a fluent explanation cannot make an outdated record current.

CHAPTER 05

What mathematical modelling adds

AI can read a business report and offer a plausible recommendation. A mathematical business model adds a different capability: it makes selected relationships and decision rules explicit enough to be applied consistently and examined when challenged.

A model can show how the same operating condition affects several outcomes, how competing activities use shared capacity, and how a proposed change compares with the current position. The value lies in disciplined reasoning that remains traceable as the number of interacting facts grows.

A decision with several consequences

Consider an illustrative manufacturer deciding whether to accept an urgent project. The project looks attractive on its own. It would also require attention from a design team already supporting other commitments, and it would depend on a production slot that may be needed elsewhere.

The sales opportunity, design workload, delivery promises, and customer relationships are individually understandable. The difficult question is their combined effect. An attractive project can be an unattractive addition at a particular time if it consumes a capability that other work needs more urgently.

A useful model represents the relevant business relationships and exposes the assumptions. Management can then compare accepting the project as proposed, changing its timing, adjusting its scope, or declining it. The model supports the comparison; leadership remains responsible for the choice.

Why the utility vector matters

A business needs to define what it means by doing well. Revenue alone cannot answer that question. A company can increase revenue while weakening cash, disappointing customers, or exhausting the people needed to deliver future work.

In the Nessiva framework, leadership defines a company-specific utility vector. The word sounds technical, but the executive idea is straightforward: keep the important outcomes visible together. Depending on the business, these may include financial health, customer value, quality, delivery reliability, and organizational resilience.

This avoids forcing every concern into a financial estimate when the estimate would be weak or misleading. A serious quality issue deserves attention even when its eventual cost is unknown. A loss of specialist capability may matter before it appears in an accounting report.

A summary view can be useful, provided leaders can still inspect its components and understand the priorities behind it. The model should help reveal a trade-off instead of concealing it inside an apparently precise total.

Explicit rules make disagreement useful

Two managers may disagree because they have different information, use different definitions, or value different outcomes. Those are distinct disagreements. A structured model helps identify which one is present.

If the disagreement concerns a factual input, the team can examine the source. If it concerns a definition, the definition can be clarified. If it concerns priorities, leadership can make a decision openly. This is more productive than arguing over competing impressions without knowing why they differ.

Objective measurement does not mean eliminating judgment. It means distinguishing observation from judgment, recording assumptions, and applying agreed rules consistently. Where evidence is limited, the model should make the limitation visible.

Scenarios prepare decisions

A model also allows leaders to ask what would change under different conditions. What if a specialist becomes unavailable? What if demand rises in one business stream? What if a review process becomes faster but production capacity stays the same?

These questions do not require the model to predict the future perfectly. They require a clear statement of the conditions being examined and an honest account of how the model behaves under those conditions. Stress tests can reveal a fragile decision even when management cannot assign a reliable probability to the adverse event.

Monte Carlo simulations extend this examination across many varied cases. More cases do not compensate for poorly chosen assumptions. Expert review remains necessary to determine whether the cases represent uncertainty that is relevant to the decision.

The role of AI in the calculation

Generative AI can help prepare data, write software, organize tests, and explain results. Where exact calculations matter, approved computational tools should execute the defined model. An AI-generated explanation should be checked against those outputs and the source evidence.

This separation is useful to executives. A coherent paragraph is not evidence that a calculation is correct. A correct calculation is not evidence that the underlying model is appropriate. Both the mathematical representation and its implementation need review.

The combined approach gives each element a clear job. Business knowledge establishes meaning. Mathematics represents selected relationships. Tested software performs calculations. AI helps build, operate, and explain the system. Management determines how the resulting evidence will influence action.

CHAPTER 06

The mathematical skeleton

A business that wants AI to support coordinated management decisions needs a common structure. We call this the mathematical skeleton. We use this image to describe the structure that gives many separate capabilities a coherent business purpose. The mathematical skeleton identifies what the business values, how selected parts of the operation relate to those outcomes, what can be measured, and which constraints a decision must respect.

An agent can prepare a report. Another can compare documents. A third can organize meetings. Their outputs become more useful to management when they refer to the same business structure and contribute to decisions that matter. Otherwise, the organization can accumulate polished activity without gaining a dependable view of performance.

Start with meaning

Suppose a team asks AI to improve efficiency. The request is incomplete. Does efficiency mean less time per document, fewer corrections, more completed projects, lower overtime, or better use of a specialist? Each interpretation can lead to a different recommendation.

The skeleton begins by making such meanings explicit. It identifies the outcome, the relevant work, the observation required, and the relationship being examined. If faster preparation is expected to improve delivery, the model should expose that connection and the conditions on which it depends.

This discipline prevents a local measure from silently becoming the objective of the whole business. A quotation team could become faster by omitting useful clarification. An inspection team could appear more productive by recording fewer issues. The business needs measures that reflect acceptable work and the consequences of shortcuts.

Connect the parts of the business

A department usually sees the work in front of it. Management needs to understand the demands shared across departments and business streams. A new sales commitment can consume design time, create procurement pressure, change installation scheduling, and affect cash at different moments.

The mathematical skeleton represents the relationships selected as important for the decision. It does not attempt to reproduce every detail of the organization. Its usefulness comes from making the consequential connections visible and calculable without overwhelming the people who use it.

In the Nessiva approach, *Virgilio* supplies this mathematical skeleton. Its utility vector states which outcomes remain in view together. Stage 1 establishes the Type A foundation across outcomes, streams and capabilities. Stage 2 deepens selected areas through Type A-B work. These elements give local AI applications a place within a company-specific management structure.

Give each tool a defined contribution

An information connection retrieves a record. A kernel supplies the business definitions and operating rules. A skill describes how to perform a task. An agent coordinates permitted steps. Approved software executes the calculations. Management reviews the resulting evidence and decides what action is justified.

These responsibilities can be implemented in different technologies. Their separation is an organizational design principle. The executive should be able to ask which component supplies the source, which applies the calculation, and which person owns the decision.

For example, an agent might locate production records and propose which entries belong to an engineering review measure. A controlled calculation can process approved entries using the defined method. A manager can inspect the exceptions and assess whether the result reflects the work. Each contribution is useful because its role is clear.

Why connection alone is insufficient

Connecting an email account, a calendar, and a spreadsheet solves an access problem. It does not establish whether a calendar entry

represents a firm commitment, whether a spreadsheet is current, or whether an email changes an approved delivery date.

Those distinctions belong in the business rules. The system needs a way to identify proposals, approvals, actual completions, and unresolved statements. If it merges them indiscriminately, access to more information may increase confusion.

Similarly, a meeting calendar may show that a specialist is free on Thursday afternoon. It cannot by itself establish the specialist's sustainable production capacity. That requires knowledge of the work, preparation, other obligations, and the standards that must be met. The mathematical structure gives operational information the meaning needed for analysis.

Governance gives the structure operating authority

A shared analytical language becomes useful across a company when the organization agrees who controls its meaning and how it informs action. Governance establishes those decisions. It specifies who approves an outcome measure, accepts a model revision, or authorizes an agent to move from preparing advice to executing a task.

This completes the operating design. The mathematics supplies a defined basis for analysis; AI helps apply it efficiently; governance makes the resulting work accountable. Cybersecurity protects the sources, methods, and permissions on which the arrangement depends. Each contribution addresses a different requirement, and none substitutes for the others.

Consider a capacity warning. A calculation may identify pressure, and an agent may assemble the relevant project information. The business still needs a responsible manager, an agreed response, and a record of what was decided. Without that connection to authority, the warning may remain an interesting observation. With it, the analysis can influence the work.

Make the structure durable and revisable

A skeleton must remain recognizable over time if the business is to compare its condition from one period to another. A delivery measure should not quietly change its definition when the reporting tool is upgraded. A shift in management priorities

should not rewrite the interpretation of every earlier decision without explanation.

Durability does not require rigidity. The business may add a product line, reorganize work, or discover that a measure is misleading. The model should then change through a recorded decision. Earlier results retain their original basis, and management is told which comparisons remain valid.

This is a valuable form of ownership. The company accumulates defined knowledge about its operation, together with a history of how that knowledge developed. The structure can outlast an individual employee, consultant, software interface, or AI model.

Use the depth that the decision requires

A simple writing task does not require a company-wide mathematical model. A standard checklist may be enough for a narrow, stable process. The need for a mathematical skeleton becomes stronger when decisions involve shared capacity, competing outcomes, repeated measurements, or consequences unfolding across time.

Begin with enough structure to support a useful decision. Add detail when the business question and available evidence justify it. The objective is a system leaders can use, challenge, and maintain. The value of mathematical sophistication is the clarity and decision support it produces.

Virgilio and *Dante* work together

Virgilio first collects and structures evidence across the organization using agreed definitions. It applies the approved business model to that evidence, producing a repeatable view of business health and preserving comparable observations over time. This includes the operating detail needed to interpret the utility-vector outcomes.

Dante then draws on that structured evidence to examine relationships. For example, a regression may investigate how a delivery outcome is associated with workload, capacity or process conditions. A time-series analysis may examine whether a change appears immediately or after a delay. These analyses help management formulate and test explanations; a correlation alone does not demonstrate a causal effect.

Both families are designed for AI integration. Their calculation core is deterministic: approved inputs, methods and configurations produce reproducible calculations. *Dante*'s statistical estimates still carry uncertainty. AI sub-processes can prepare evidence, propose classifications, flag inconsistencies, coordinate approved analysis and draft explanations. These supporting tasks make the mathematics more practical and useful without transferring control of its rules to generated text.

The operating sequence is structured collection through *Virgilio*, validated analysis through *Dante*, and management review of the findings. Accepted insights can guide action and a documented refinement of measures or assumptions. Earlier observations retain their original definitions so that model development does not silently rewrite the historical record.

CHAPTER 07

Progressive model development

A company-wide model can become unmanageable if it tries to describe every detail before it becomes useful. A very simple model can also fail if it overlooks the relationships that matter. Nessiva addresses this through progressive development.

Within *Virgilio*, Stage 1 establishes the Type A structure and Stage 2 adds depth through Type A-B work where better evidence or diagnosis is needed. Later development can extend the measurement structure and connect *Dante*'s analytical models when sufficient comparable evidence is available. The scope follows the business question; the two model families work together.

Type A establishes the management view

The first task is to agree on the outcomes the business values and the principal streams and capabilities through which it operates. This gives leadership a common language for discussing what is happening and where its effects are felt.

For a manufacturer, streams might distinguish categories of projects or products. Capabilities might include commercial work, engineering, production, installation, and service. A retail organization could use product categories or sales channels as streams and purchasing, merchandising, fulfilment, and customer support as capabilities.

The categories must reflect the business. They are not an organizational chart copied into a template. A capability may support several streams, and one project may require several capabilities at the same time. Those shared demands are often important to management decisions.

Type A connects this structure to a repeatable view of performance, capacity, and risk. It can use defined expert assessments where more direct measurements are not yet available, provided their basis and limitations remain visible. Its

purpose is to establish useful orientation and identify where deeper investigation is worth the effort.

Manufacturing implementation experience

In a manufacturing engagement, the Type A foundation has already been completed. The delivered work includes a company-specific mathematical model, a working management workbook, supporting documentation, and an AI kernel prepared around the model.

The completed foundation gives management a working analytical tool and a common language for further investigation. Its practical value can then be evaluated through use, while the next stage develops more detailed operating evidence.

The continuing programme is designed to deepen selected operational measurements while developing targeted AI applications and their governance. This distinction matters: completed infrastructure, work in development, and realized business benefits are different milestones.

Connect the view to operating evidence

Once leadership can see the business structure, it can choose a capability or process that deserves closer attention. Type A-B investigates that area and develops a more precise connection between operating evidence and the broader management view.

Suppose engineering is the selected area. The work begins by understanding what the department is responsible for, what constitutes acceptable output, and which records can show whether those requirements are met. It then examines workload, resource use, and operational strain in a way that suits the work.

The sequence is important. The team first defines what needs to be understood, then identifies meaningful evidence, then develops the measurement logic. A convenient formula should not determine the meaning of the business question.

This work benefits from direct involvement by the manager and the people doing the tasks. A general description of manufacturing is a starting hypothesis. Local staff establish whether it reflects the operation being modelled.

Capacity requires more than recorded hours

An employee's recorded hours show resource consumption. They do not establish how much work can be sustained at an acceptable standard. A team may complete its work through repeated overtime, rushed checks, or postponing activities that will become urgent later.

Equally, waiting time does not automatically establish excess capacity. A person may be waiting for information while a queue of dependent work continues to grow. Different skills can be constrained at different times within the same department.

A useful investigation therefore considers the period, the work mix, completion requirements, and evidence of pressure. It may draw on queues, overtime records, rework, delivery performance, and documented observations. The aim is a defensible view of what the operation can sustain.

Development over time

Repeated measurements allow management to examine whether conditions are improving, becoming more variable, or responding differently to changes in demand. This history may support time-series analysis, regression, and forecasting.

Such methods require sufficient, comparable evidence. A new measure with only a few observations may be useful for discussion while remaining unsuitable for a reliable forecast. Changes in definitions or operations must also be considered when interpreting the history.

Stage	Management purpose	Typical output
Type A	Establish the company view	Defined outcomes and an initial model of operating relationships
Type A-B	Investigate a priority area	More precise measures connected to operating evidence
Later development	Understand change and support further decisions	Repeated observations, validated extensions, and suitable scenario or forecasting tools

The benefit of this progression is controlled investment. Management can learn from a useful foundation, deepen the areas

that justify attention, and expand the model when the evidence and business need support it.

CHAPTER 08

A business model that lives over time

A single measurement describes a position. Repeated, comparable measurements describe movement. That movement is often where the most useful management information appears: a queue is growing, a recovery is losing momentum, variability is increasing, or an improvement is holding under greater demand.

Maintaining the mathematical skeleton over time creates a record of business conditions and management decisions. Consistent definitions help distinguish a real change from a change in reporting.

Give every observation a time meaning

A number needs more than a label. Management must know whether it describes a particular moment, an activity during a period, or an estimate of what will happen later. Outstanding jobs at month-end, jobs completed during the month, and jobs expected next month answer different questions.

These observations should remain distinct. Adding several month-end backlog totals does not establish how many separate jobs entered the business. The same unfinished job may appear in more than one snapshot. Equally, a monthly completion total says little about the work still waiting unless the records connect the two.

The reporting period also matters. An unusually short working month may change activity totals without indicating a deterioration in performance. Where useful, the analysis should account for relevant differences in working time and work mix. The adjustment must be defined rather than improvised to explain an inconvenient result.

Preserve what was known when

A useful record distinguishes when something happened from when the organization learned about it. A defect found today may concern work completed earlier. A supplier may report a delay after management has already committed to a schedule. A corrected invoice may change the apparent position of a previous period.

Preserving these distinctions helps leaders assess their decisions fairly. The question is often whether the action was reasonable given the evidence available at the time. A report rebuilt entirely from today's corrected information may conceal that earlier uncertainty.

When a record is corrected, retain an explanation of the correction and a connection to the earlier version. Where management reports are restated, identify that change. This gives the business both a current view and an understandable history.

Follow several outcomes together

A business may improve delivery while using more overtime. It may improve cash by reducing stock while increasing exposure to supply disruption. Looking at each measure separately can obscure the connection.

Following the utility vector over time keeps these movements visible together. Management can examine whether an apparent gain is sustained by acceptable operating conditions or by drawing down something the business will later need. The purpose is to understand the pattern of gains and sacrifices.

An illustrative manufacturer might track release quality, delivery reliability, and engineering strain across successive periods. If release becomes faster while correction work increases, leadership has a reason to examine how the time saving was achieved. If both quality and delivery improve under comparable demand, the evidence supports a more promising interpretation.

Recognize timing, repetition, and delay

Time-series methods examine observations in their temporal order, including patterns such as trend, seasonal variation, and dependence between successive periods. This is a long-established analytical discipline. [4]

The executive application begins with ordinary questions. Is the change persistent? Does it recur at a similar point in the business cycle? Does one measure tend to move before another? Has the relationship changed since a new process was introduced?

A delayed effect can be particularly important. Better document preparation may reduce correction work later. A sales campaign may create delivery pressure only after orders reach a constrained capability. A staffing change may take time to affect throughput because training temporarily consumes experienced staff attention.

The analysis should reflect the time needed for the proposed mechanism to operate. Evaluating a change too early can miss its benefit. Waiting too long without interim evidence can allow a failing approach to continue. Management needs a review schedule suited to the actual business process.

Build a history fit for analysis

Longer records can support better questions and more demanding analysis, provided their quality and relevance are sufficient. A large archive of inconsistent measures can be less useful than a shorter, well-defined series. Missing observations also need to remain visible; absence of a record does not establish zero activity or zero risk.

The past may cease to represent present conditions. A new product mix, major customer, production method, or purchasing policy can change the relationships on which earlier analysis relied. The model must identify such changes and reconsider what the older data can support.

Forecasting should therefore develop when there is a justified business use and enough suitable evidence. Early records may support careful monitoring before they support dependable prediction. This is a sensible progression in model maturity.

Make continuity part of the operating routine

AI can reduce the recurring effort of collecting approved records, applying established checks, preparing comparisons, and drafting explanations. The continuity comes from the maintained records and methods used by that workflow. It should not depend on someone remembering to paste last month's answer into a new conversation.

CHAPTER 09

Traces and reliable calculations

When a business asks why performance changed, it needs more than the latest chart. It needs a trace: a preserved path through relevant observations, calculations, decisions, and subsequent events. That path allows a manager to examine what changed, on what evidence, and with what consequence.

The record connects received information, approved methods, outputs, authorized actions, and later business events. It concerns observable work, not a language model's hidden internal reasoning.

Keep a trace that serves a question

There is little value in retaining every possible detail without knowing what someone will need to examine. A trace should be proportionate to the decision and the information sensitivity. A recurring management measure needs enough detail to explain the number and reproduce the relevant calculation when required.

For example, the record may identify the reporting period, the source records used, unresolved data issues, the version of the calculation, and the person who accepted the report. A consequential action needs an additional record of the approval and what was actually executed.

The same history can support several views: an executive explanation with supporting evidence, individual exceptions for a process owner, and calculation versions and test results for a technical reviewer.

What a carefully written kernel contributes

A kernel can define the approved terms, required sources, calculation procedure, exception rules, and expected form of the result. It can instruct an agent to distinguish a missing input from a zero, preserve dates, cite source records, and refer unresolved ambiguity to a named owner.

This makes recurring work easier to standardize. The business does not need to restate its operating meaning from the beginning on every occasion. A skill can then describe a specific recurring task, such as preparing the monthly capability review using the current approved kernel and calculation package.

Instructions alone do not enforce reliable execution. Important requirements should also be reflected in software checks, access controls, and acceptance criteria. If a reporting period is missing, the workflow can reject the input instead of merely asking the language model to notice the omission.

Let approved software perform exact calculations

Where mathematical formulas govern an output, the formulas should be implemented in a controlled calculation tool and checked against independently established expectations. AI can prepare inputs, invoke that tool, and explain its results. The explanation should refer to the recorded output rather than invent or casually recompute the answer.

For calculations intended to be deterministic, the same approved inputs and method should produce the same result. Where a method uses randomness, its settings and the information needed to reproduce the run should be recorded. Neither form of reproducibility proves that the business assumptions are sound; it establishes that the computational process can be examined.

A calculation can be exact and still answer the wrong question. If a measure treats waiting time as productive capacity, correct arithmetic will preserve that conceptual error. Quantitative expertise is needed to establish the representation, and operating expertise is needed to test whether it describes the business adequately.

Preserve the chain across repeated runs

A recurring workflow needs to recognize which records have already been processed, which have changed, and which remain incomplete. A second run should not quietly count the same event twice or repeat an already completed external action. A failed run should not be mistaken for a period with no exceptions.

The design should therefore include durable storage, clear record identifiers, checks on completeness, and visible run status.

Source records may remain in existing business systems. The analytical workflow can preserve references and approved snapshots where required, subject to the organization's information rules.

AI helps interpret and coordinate the work; controlled software and records preserve continuity. A strong kernel connects these elements, and monitoring checks their operation.

Connect decisions to later outcomes

Suppose management changes a review process to reduce late corrections. The decision record should state the intended effect, the part of the process changed, the implementation date, and the evidence that would justify continuation. Subsequent observations can then be examined against that expectation.

If performance improves, the trace helps ask whether the improvement followed the intended mechanism. Did discrepancies move earlier in the process? Did correction work fall? Did reviewers have more time for difficult cases? These questions are stronger than simply observing that the new system was present during a good month.

Other changes may also explain the result. Work could have become easier, demand could have fallen, or a new employee could have joined. Where practical, comparisons with unaffected work or a staged introduction can strengthen the evaluation. Where they are impractical, the business should describe the uncertainty honestly.

Test predictions using the past fairly

When a forecasting method is introduced, it can be evaluated on historical periods by recreating what could have been known at each forecast date. Later information must not slip into the earlier prediction. The original forecast should be retained so that it can be compared with the actual outcome.

Compare the method with a simple, relevant alternative. A complex forecast should earn its additional cost through better performance for the intended decision. Management should also examine the consequences of error: a forecast that is adequate for staffing discussion may be inadequate for a firm delivery promise.

CHAPTER 10

Where AI improves everyday work

The most useful first applications often involve information that people already handle repeatedly. A team reads documents, compares versions, gathers evidence, checks requirements, or prepares a standard output. These activities may offer a manageable place to test AI because the work and its result can be inspected.

The following illustrative applications show how the combined approach can improve ordinary work. Each explains the source of the potential benefit and the operating conditions needed to realize it.

Fabrication sheets and production documents

A fabrication sheet tells people what needs to be made and includes the information needed for production. Problems can arise when it omits a required field, refers to an outdated document, or conflicts with an approved specification.

A suitable AI workflow can help assemble the relevant documents, check defined information, and prepare a list of possible discrepancies. For each finding, the output should identify the location in the source and explain what appears inconsistent. Unreadable or missing material should be reported as a limitation.

The reviewer can then focus on the flagged issues while retaining responsibility for the complete review required by the process. Potential value comes from reduced search and comparison effort, earlier discovery of omissions, and a clearer record of the checks performed.

For example, the workflow might identify that a production sheet cites one revision while the accompanying package contains another. The system should flag the mismatch. The person authorized to establish the controlling revision must resolve it before release.

Drawings and revision comparisons

AI can assist with extracting readable annotations, organizing revision differences, and comparing stated requirements across documents. A drawing review may also require specialist software to examine dimensions, geometry, or relationships that are not reliably represented in extracted text.

This distinction defines the scope of a useful pilot. If the task is to compare readable revision notes, the system should be tested for that task. It should not present the result as a complete engineering verification of the drawing.

The possible benefit is a better-prepared review package: changes grouped for attention, source locations preserved, and unresolved items made visible. Qualified engineering personnel remain responsible for technical interpretation and approval.

Contracts and commercial commitments

A contract review workflow can help extract obligations, identify differences from an approved commercial template, and assemble questions for the responsible reviewer. It may highlight a delivery term that differs between a quotation and the proposed contract, or an obligation that is not reflected in the operating plan.

The business value is that relevant commitments become easier to locate and discuss before acceptance. A useful output cites the actual clause, identifies the comparison being made, and avoids treating a summary as the agreement itself.

The workflow should preserve the distinction between commercial preparation and legal judgment. It can help the business ask better questions and prepare for professional review. It cannot confer authority to accept terms or establish their legal effect simply by producing an answer.

Quotations and customer communication

Preparing a quotation often requires collecting known information, applying current templates, and identifying what still needs clarification. AI can help with those tasks and prepare a draft that is easier for the responsible employee to review.

The workflow should use approved pricing sources and commercial rules. It should make missing information visible

rather than silently filling a gap. A quotation can look polished while containing an unsupported delivery date or an assumption the business would never accept.

Faster preparation can improve responsiveness and leave more time for understanding the customer's needs. The final benefit depends on maintaining accuracy and preserving the company's commercial authority.

Retail purchasing and inventory decisions

Our team's retail business experience informs the following illustration of how AI and mathematical modelling can support purchasing decisions.

Consider a retailer with both excess stock and recurring shortages. An AI assistant can help organize supplier communications, compare product descriptions, and prepare an explanation of inventory exceptions. A mathematical model can examine the relationships among demand, stock, replenishment timing, cash commitments, and the service level the retailer wants to provide.

The distinction is useful. Reading supplier messages efficiently does not determine the appropriate stock policy. A stock model does not automatically interpret an unstructured supplier message. Connecting the two can give the buyer more timely information and a clearer basis for a purchasing decision.

An initial application could prepare a review list of items whose recorded stock or supply conditions deserve attention. Purchasing authority remains with the buyer until any wider automation has been specifically tested and approved. Possible benefits include less administrative effort, more consistent exception review, and purchasing decisions that better reflect the business's priorities.

Service delivery and internal knowledge

A service organization may depend on specialists who repeatedly answer questions that are already addressed in approved records. An AI assistant can help employees retrieve the relevant guidance, summarize it, and identify questions that need escalation.

A useful response includes the source and its status. If the approved procedure is silent, the assistant should say so. If two

records conflict, it should identify the conflict rather than choose a convenient answer without authority.

This can support onboarding, reduce repeated interruptions, and make established knowledge easier to use. Specialist time becomes available for cases that require interpretation, client judgment, or an exception to normal practice.

Management reporting

AI can help assemble recurring reports from approved operating records, compare current information with an earlier period, and prepare a readable account of changes. It can also help leaders question a controlled record of an ongoing engagement: what was planned, what was delivered, and what remains dependent on another action.

The source set must distinguish plans from completed work. A report dated last month does not become a factual completion record merely because an assistant can retrieve it. The workflow must preserve status, dates, and unresolved questions.

This application connects naturally to the mathematical model. The model supplies a defined analytical view; the assistant helps people understand and question that view. A manager should be able to ask why a result changed and receive an explanation grounded in the relevant inputs and records.

Select applications by the work they improve

Across these examples, the useful pattern is consistent. Define a bounded task, identify authoritative information, determine what the system should produce, and place review where it can affect the result. Then measure whether the whole workflow improves.

The first choice should reflect business value, available evidence, and the ability to detect and manage error. A technically simple application can be commercially important if it removes a recurring obstacle in a constrained part of the business.

CHAPTER 11

Make AI part of a controlled workflow

A personal experiment with an AI assistant can reveal an opportunity. A business workflow requires that opportunity to survive repeated use by different people, with changing documents and real consequences.

The design begins with the work. What triggers the task? Which information is authoritative? What output is required? Who receives it? Which decisions remain with a person? Only after these questions are understood should the team choose whether an assistant, an agent, conventional automation, or a combination is appropriate.

Why context needs an owner

A general AI system does not automatically know the company's current definitions, approved priorities, or authority structure. Even a capable system can give an unsuitable answer when it lacks the right context.

A company configuration supplies relevant approved material. This may include policies, terminology, templates, model descriptions, and factual records. Someone must own that material, keep it current, and decide which users may access it.

A model-specific kernel can help preserve the analytical language and operating rules. It can instruct the assistant to separate evidence from assumptions, identify missing information, and explain outputs within the approved framework. Reusable skills can then guide recurring tasks within that environment.

This is how expertise becomes easier to apply repeatedly. The organization records part of its working knowledge in a form that supports consistent execution and review. The arrangement still needs tests and enforceable access controls; written instructions alone cannot guarantee compliance.

Define the authority of an agent

An assistant can help a person inspect a document or prepare a report. An agent may gather documents, invoke a calculation tool, prepare findings, and route an item through several steps. That additional capability is useful when it reduces repeated coordination work.

It also means the organization must define the agent's permitted actions. Reading a record, drafting an amendment, and applying that amendment are separate permissions. A system that is useful for preparation does not automatically need authority to change production records or send a commercial commitment.

For example, an agent might prepare a list of potential document discrepancies and create an internal review item. The responsible reviewer decides whether a discrepancy is real and authorizes the next step. The review decision should be recorded where the actual process relies on it.

Use the simplest system that does the job

Some tasks have stable rules and structured inputs. Conventional software may perform them reliably with little need for generative AI. Other tasks involve varied language, incomplete documents, or questions that benefit from conversational explanation.

A useful workflow can combine both. AI helps interpret and organize the information; defined software performs exact calculations or applies stable checks. This reduces the burden placed on any one component and makes failures easier to investigate.

Adding more agents does not necessarily improve the result. Several systems working from the same incorrect source can repeat the same mistake. Independence of review must be designed through appropriate evidence, methods, or qualified people.

Make uncertainty actionable

A controlled workflow should have a useful response to missing, conflicting, or unreadable information. It may ask for a clearer document, flag a contradiction, or stop the affected step for review. This allows the business to act on uncertainty instead of hiding it.

The output should help the reviewer answer three practical questions: what was examined, what was found, and what still needs attention. These questions keep the review connected to evidence and make it easier to determine whether the system has stayed within scope.

Maintain the workflow after launch

Documents change. Software is updated. Staff responsibilities shift. A supplier starts using a different format. Each can affect a previously useful workflow.

The organization therefore needs an owner, a record of significant changes, and a way to retest the cases that matter. It also needs a workable fallback when the AI service is unavailable or the output cannot be trusted. The fallback may be the existing manual process with a clear route for urgent cases.

Integration is complete enough to operate when people understand the process, the controls work in practice, and someone is responsible for keeping it useful. Continuing maintenance is part of the business case.

CHAPTER 12

Connect the business through ChatGPT

Many companies already use AI individually. One person drafts correspondence, another studies a spreadsheet, and another prepares meeting notes. A business workspace creates an opportunity to organize this activity around shared responsibilities and approved resources.

ChatGPT Business is one possible environment for that work. OpenAI describes it as providing a dedicated workspace with essential administration controls and business data excluded from model training by default. Specific features, model access, and usage allowances depend on the applicable plan and configuration. More extensive enterprise controls should not be assumed to be included in every Business subscription. [5]

The executive decision concerns the operating arrangement built around the account. Who owns shared resources? Which systems may be connected? What may an assistant read or change? Where are accepted outputs retained? A subscription makes capabilities available; management gives those capabilities an organizational purpose.

Connect information where it already belongs

Plugins can provide access to information and supported actions in external services. Current OpenAI documentation describes integrations for tools such as Gmail, Google Drive, and Slack. The connection must be available and authenticated, and its actual tools determine what the workflow can do. [6]

A business can use this pattern to reduce repeated copying between systems. An agent might retrieve the approved project file, inspect the relevant operating record, and prepare a report with references to both. The original systems can continue to hold their respective authoritative records.

Connecting a service does not make every document available to every colleague. Each workflow should respect the authenticated user's access and the organization's policies. The intended outcome should determine the access needed. A management summary may require selected aggregates without requiring every participant to see individual personnel records.

Calendars become useful operating context

OpenAI's current training materials describe Google Calendar capabilities for checking availability and creating events, alongside email and other integrations. Available plugins and actions vary by plan and workspace. [6]

The business use extends beyond finding a meeting time. Consider an illustrative manufacturing review that requires a commercial manager, an engineering reviewer, and a production planner before a proposed release date. With appropriate connections, an AI workflow could help coordinate the discussion and assemble the approved material those people need.

The skeleton gives that coordination its meaning. It identifies the release condition, the required roles, the evidence to review, and the consequence of a delay. A free calendar slot is useful because it serves a defined step in the operating process.

The authority to read availability should be distinguished from the authority to create or alter events. Attendee identities, time zones, and the status of a proposed commitment need verification. A scheduling aid becomes dependable when it handles these ordinary details consistently and records what was actually agreed.

Projects and custom GPTs serve different purposes

A project can hold related work and context. A custom GPT can provide a reusable assistant configuration for a defined purpose. Sharing a GPT can make an approved assistant accessible to colleagues, but it should not be treated as a universal shared database or an automatic connection among all their conversations.

OpenAI's administration documentation describes controls for GPT sharing and ownership, and the use of connected apps or custom actions under applicable permissions. It states that a GPT can use connected apps or custom actions, but cannot use both

together in the same GPT. The available controls should be checked for the actual workspace. [7]

An illustrative company might define an assistant for approved policy questions and a separate analytical workflow for management reporting. Both can refer to consistent business definitions where their supported configuration allows it. Their information access and responsibilities can still differ.

The company should maintain an authoritative version of shared rules and an explicit process for updating each implementation. A revision to one GPT's instructions should not be assumed to update every project, skill, or agent. Configuration has a lifecycle just as an operating procedure does.

Repeat work on a deliberate schedule

Scheduled tasks can run recurring work, and supported app events can trigger tasks on eligible plans. OpenAI documents the use of accessible files, connected tools, skills, and plugins within supported scheduled workflows. Local work may require an active computer and application; a cloud workflow has different execution conditions. [8]

The executive application is a regular operating rhythm. A workflow could prepare a weekly exception review after the relevant records are available. Another could assemble a monthly comparison using an approved reporting period. The schedule should follow the availability and meaning of the evidence.

An event trigger is useful when action depends on something happening, such as a relevant incoming update. It does not replace the procedure for deciding whether the update is authoritative or important. A new message may contain a proposal, a duplicate, or incomplete information. The workflow still needs interpretation and controls.

Repeated tasks also need a visible completion state. Someone should know if the report failed, a connection expired, or the required records were absent. A quiet system is only reassuring when its monitoring can distinguish normal operation from a failure to run.

Keep connection separate from commitment

An interconnected environment can prepare many actions: proposed meetings, customer replies, purchasing requests, and updated reports. The business should specify which actions require a person's approval and which may proceed under a standing rule.

For example, collecting source information and drafting an internal report may be delegated within a defined scope. Changing a customer commitment may require a named manager. The rule should reflect the business consequence and be enforced by the available technical controls.

This arrangement allows useful autonomy. Staff spend less time assembling routine material, while responsibility remains visible at the points where business commitments change. It also makes performance easier to assess because the system's delegated task has a clear beginning and end.

Start with one connected decision

A first connected workflow need not include every application. Select a decision that repeatedly requires information from a few places. Define its evidence, owner, calculation requirements, and expected output. Connect the necessary sources and test the complete sequence.

The result may reveal that a missing record definition matters more than another integration. It may also reveal that a well-maintained source can support several later applications. These are useful findings because they guide investment towards business value.

Nessiva can help design this operating structure: the mathematical meaning, AI configuration, information connections, and governance around them. The purpose is to make the company's existing tools contribute to a coherent management process.

CHAPTER 13

Agents and kernels working together

Once a business has several useful AI applications, the next question is how they should cooperate. Cooperation requires more than asking several agents to discuss a problem. Each participant needs a defined task, a recognized input, an expected output, and rules for handling disagreement or incomplete work.

The common mathematical skeleton gives these tasks a shared frame of reference. A document-review workflow and a management-reporting workflow may work at different levels, yet both should understand the same definition of an approved release or completed corrective action.

Three kernel designs

An autonomous kernel guides a workflow through permitted steps with limited intervention for each routine action. It may gather approved records, check whether an evidence package is complete and prepare a recurring report. Its mandate defines the tools it may use, the actions it may take, stopping conditions and the cases it must escalate. Autonomous operation remains bounded by the organization's authority and access rules.

A deterministic kernel runs approved calculations according to explicit rules. For the same approved inputs and configuration, it produces the same result. This design supports repeatable *Virgilio* measurements and reproducible *Dante* analyses. AI can help assemble the input package and explain the output; tested calculation software executes the mathematical method.

A judgment-embedded kernel includes AI sub-processes where interpretation is necessary. For example, AI may propose how a manager's narrative maps to a defined performance category, identify an ambiguous record or draft a reason for an exception. The process records the evidence, proposed interpretation,

uncertainty and review status. Accepted values then enter the deterministic calculation stage.

These designs can cooperate within one workflow. An autonomous process can coordinate collection, a judgment-embedded step can propose a classification, and a deterministic kernel can calculate the result after the required checks. AI-assisted judgments may vary, so the reproducibility claim applies to the approved calculation stage. Material interpretations require review under the agreed rules.

This is how AI sub-processes strengthen a deterministic model: they reduce the effort of preparing and interpreting evidence while keeping the mathematical relationships explicit and controlled. For *Virgilio* and *Dante*, the resulting capability combines computational discipline with practical assistance in the work surrounding the calculation.

Use specialization where it helps

An illustrative operating arrangement might assign one workflow to collecting and checking source records, another to running approved calculations, and another to drafting the management explanation. A designated reviewer then examines the material before it informs a consequential decision.

These roles need not all be separate AI agents. A conventional program may be the best calculation component. A simple form may be sufficient for a manager's approval. The architecture should make the work easier to understand and maintain.

Separate agents can be useful when their tasks can be clearly bounded or completed independently. They can also create coordination costs and inconsistent interpretations. Adding agents is justified when the improvement in the complete process exceeds that additional burden.

Define the handoff

A handoff is a statement of what the next step may rely on. A document-review result might identify the revision examined, the discrepancies found, the evidence locations, and the issues still awaiting judgment. It should not describe unresolved work as approved simply because the review task has finished.

The calculation stage needs inputs whose meaning and status are clear. The reporting stage needs approved results and a way to distinguish an observed result from an assumption. Management needs to know whether the full evidence package is complete enough for the decision.

This clarity prevents a small ambiguity from growing as information passes through several components. A tentative interpretation should not become an unquestioned fact merely because another agent repeats it in more confident language.

Share rules through controlled versions

A common kernel can express the business language used across workflows. Individual skills can specify their procedures and outputs. Access permissions and executable controls then define what each implementation may actually do.

The approved version matters. If one workflow uses an old completion definition while another uses a revised one, their reports may appear comparable when they are not. Each consequential output should identify the relevant analytical basis, and changes should be distributed through an explicit update process.

The same principle applies when a stronger AI model becomes available. Management can preserve the business definitions while evaluating the new model's role. The company improves a component without casually changing the meaning of the entire system.

Arrange independent checking

Several agents can make the same mistake if they rely on the same incomplete source or ambiguous instruction. Agreement among them is therefore limited evidence. A second fluent explanation is not equivalent to an independent examination of the calculation or source record.

Useful checking introduces a different basis for assessment: a tested rule, a known reference case, direct source verification, or a competent person's review. The check should address the failure that matters. Repeating the original prompt may add little protection against a shared misunderstanding.

This is especially important for drawing interpretation, contractual obligations, and business measures. The appropriate reviewer needs the expertise to assess the consequential point, together with enough evidence to do so efficiently.

Make recovery part of the design

A connected process can fail partway through. A source service may be unavailable after information has been collected, or an external action may succeed even though its acknowledgement is lost. The workflow needs to determine its actual state before repeating the action.

For executives, the requirement is straightforward: staff must be able to see what completed, what remains pending, and what needs intervention. The system should preserve the last accepted state and offer a practical recovery procedure. Continuity becomes more important as more work depends on the connection.

Build a business capability that can develop

Over successive periods, a coordinated system can become more useful through better source records, clearer handoffs, improved tests, and carefully selected new AI capabilities. The mathematical skeleton supplies continuity; the surrounding tools can develop as business needs change.

This is the strategic opportunity for a smaller organization. It can acquire a disciplined analytical and operating capability in stages, supported by a compact expert team. It can improve practical work now while developing the records and structure needed for more informed decisions later.

The following chapters address the governance, leadership, and measurement needed to sustain that development. These responsibilities are part of the capability itself. They help ensure that increasing technical power produces results the business can understand and trust.

CHAPTER 14

Governance makes the system accountable

Mathematics gives a business an explicit structure for examining its decisions. AI helps develop that structure, apply it repeatedly, and connect it with everyday work. Governance gives the combined system its operating authority: who may use it, what it may influence, how its conclusions are challenged, and who remains responsible for the consequences.

This is how the analytical structure becomes a business capability. A model may be well designed and its software accurate, yet the organization still needs to decide who approves a change in priorities or acts on an emerging capacity problem. An agent may prepare an excellent recommendation, yet its authority to make a commitment must come from the business.

Governance therefore belongs in the design from the beginning. It helps a company use powerful tools confidently within known boundaries, while cybersecurity protects the information and connections on which that confidence depends. Both must develop as the system changes.

Establish responsibility before delegating work

Management sets the outcomes the organization values and accepts material trade-offs. A process owner defines acceptable operating performance. Analytical and technical specialists establish and test the relevant methods. Users follow the approved workflow and report problems. These responsibilities should be clear even when a small team holds several roles.

An executive sponsor also needs to resolve competing demands. If sales wants faster approval and engineering needs additional evidence, the AI system should not silently decide which requirement takes precedence. It can make the disagreement visible and assemble the information needed for an authorized decision.

A named owner should be responsible for each consequential model or AI workflow throughout its life. Ownership includes keeping its purpose current, arranging reviews, and ensuring that someone can suspend or replace it. A system that outlives its original developer should retain an accountable business owner.

Match oversight to the consequence

Preparing an internal draft, recommending a purchasing decision, and changing a customer commitment deserve different levels of oversight. Management should classify a use by the information it handles, the actions it can take, and the damage a mistake could cause.

A limited drafting aid may need a straightforward review procedure. A workflow affecting delivery promises or access to sensitive records requires stronger evidence, tighter authority, and a defined response when it behaves unexpectedly. The classification should change when the scope or permissions change.

Keep a proportionate register of active uses. For each one, record its purpose, owner, approved sources, permitted actions, review requirements, and current status. This gives management a practical view of what the business is relying on and helps prevent forgotten experiments from becoming unofficial production systems.

Protect the meaning of the mathematical model

Model governance concerns the representation of the business as well as its software. A changed priority, definition, or parameter can alter a result even when operating conditions have remained the same. Earlier comparisons can become misleading if that change is hidden.

A proposed revision should identify the business reason, expected effect, supporting evidence, and authority required. Suitable tests should examine both ordinary cases and situations in which the change could produce an unwanted result. Once accepted, the revision receives a recorded effective date and a clear relationship to earlier versions.

The same care applies to the utility vector. Leadership may legitimately change its priorities, but the model should show that

decision openly. It should also preserve important constraints so that an attractive aggregate result cannot conceal an unacceptable quality or security outcome.

Staff need a route for challenging the model. If its output conflicts with credible operating evidence, the response should be to examine the data, assumptions, implementation, and intended use. A model earns trust by surviving informed examination and being corrected when the evidence requires it.

Turn policy into operating controls

An AI policy should help employees make ordinary decisions: which tools they may use, what information they may provide, what they must verify, and when they must ask for approval. A model policy should identify who controls the analytical definitions and how changes are reviewed.

The corresponding controls must work in the actual environment. A rule that requires approval before sending a customer message is incomplete if an agent's connected tools can bypass that approval. A rule restricting sensitive data needs appropriate permissions in the source and destination systems.

Policy area	Executive purpose	Example of an operating control
Information use	Preserve approved uses of data	Access restricted to the records required for the task
Business commitments	Preserve accountable decisions	An authorized person approves a consequential action
Model changes	Keep results interpretable	Significant revisions are tested and recorded
Feedback and challenge	Detect unsuitable conclusions	Staff can refer concerns to a named owner
Incident response	Limit harm and restore work	The workflow can be suspended and a fallback used

A temporary exception also needs an owner, reason, scope, and expiry or review date. Otherwise, a workaround can quietly become the permanent process. Exceptions should provide useful evidence about where the current design may need improvement.

Make human review a real capability

Review is meaningful when the reviewer understands the issue, can inspect the evidence, has enough time, and can refuse or change the proposed action. A stream of unexplained approvals can create the appearance of control while encouraging people to accept outputs automatically.

Design the review around the consequential uncertainty. A commercial manager may assess whether a proposed commitment is acceptable; an engineer may assess a drawing discrepancy; a quantitative specialist may examine a measure's validity. The information shown to each should support that particular judgment.

For more consequential work, the person checking the method should have enough independence to challenge its designer. Smaller businesses can arrange focused independent reviews without maintaining a separate department. The scope and limits of the review should be recorded.

Follow one decision through the whole system

Consider an illustrative manufacturer assessing an urgent order. The mathematical model compares the proposed work with shared capacity and existing commitments. AI gathers approved records, runs the authorized analysis, and prepares the comparison. Governance determines who may accept the trade-off and what further evidence is required. Security controls restrict the data and actions available during the process.

Management might approve the order with a revised delivery date, subject to a specific engineering review. That condition must remain visible in the operating record. An agent preparing the subsequent customer message should use the approved decision, not infer acceptance from a positive draft recommendation.

Later, the business can compare the expected and actual consequences. The trace connects the information, calculation, decision, and result. Governance makes that history usable by establishing which records are authoritative and who acts when the result departs materially from expectation.

Review the system throughout its life

Before wider use, leadership needs evidence that the workflow is useful, the controls function, and its owner can maintain it. During operation, reviews should consider performance, unresolved exceptions, changes in access, and emerging risks. A retired workflow should lose its unnecessary permissions while required records remain available under the retention policy.

The NIST AI Risk Management Framework and its Generative Artificial Intelligence Profile offer public reference points for managing AI risk across design, use, and evaluation. They support a structured approach; applying a framework does not itself certify a system or establish compliance. [9]

Nessiva brings this governance work together with quantitative development and AI integration. Our team helps turn management expectations into practical policies, responsibilities, controls, and reporting. The value is a system that the business can use, question, and improve with clear accountability.

CHAPTER 15

Cybersecurity for connected AI

Connecting AI to business systems increases what it can accomplish and what it can affect. An assistant that reads a public document has a different exposure from an agent that can retrieve customer files, alter a schedule, or initiate an external action. The security design should follow these actual capabilities.

For executives, the objective is to protect information, preserve the integrity of decisions, and keep important work available. A breach can expose confidential material, but it can also corrupt a calculation, redirect an instruction, or interrupt a process at a critical moment. Cybersecurity protects the mathematical skeleton and the work built around it.

Understand the information path

Begin with a practical description of where information comes from, where it is processed, what is retained, and where outputs go. Include the AI service, connected applications, company storage, and any third parties involved. Identify which records are authoritative and which are working copies.

This helps management ask specific questions. Does a task need complete customer files or a limited extract? Does a draft contain information that should not reach its intended audience? Will a sensitive input be copied into a shared report or retained longer than required?

Useful traces must also respect these boundaries. A record of an action may need an identifier, date, and approval without reproducing the entire confidential document. Retention and access should reflect the business purpose and applicable obligations. Recording more information is not automatically better oversight.

Restrict identities and permissions

Each connection operates through an identity with particular rights. Those rights should be appropriate to the task, reviewed when responsibilities change, and removed when no longer needed. Strong sign-in controls and careful handling of credentials remain important when an AI system uses the account.

The principle of least privilege means granting only the access required. A workflow that prepares a comparison may need to read approved records and write a draft into a designated location. It need not have authority to delete source files, change access rights, or send the result externally.

Permissions should be enforced by the connected systems and tools. A written instruction to an agent is useful guidance, but the business should not rely on the agent alone to decide whether it has authority. Sensitive credentials belong in controlled credential mechanisms, not in reference documents or conversational instructions.

Treat external content as evidence

AI can encounter instructions hidden or embedded in material it is supposed to examine. OWASP describes this as a prompt-injection risk, including indirect attempts through documents and websites. Additional reference material or model customization does not reliably remove this exposure. [10]

An illustrative procurement workflow may receive a document that attempts to redirect the assistant's task. The document should remain evidence to assess, without gaining authority over company instructions or access controls. Clear separation of trusted rules from external material, restricted tools, validation, and appropriate approval requirements help limit the consequences. [10]

The executive requirement is to test how the complete workflow handles untrusted material. A persuasive statement within a document should not be enough to change a payment destination, expose another customer's information, or authorize a business commitment.

Limit the consequences of an agent error

OWASP's excessive-agency guidance identifies risks created by unnecessary functions, permissions, or autonomy. A capable agent can cause damage when its tools allow more than the intended task requires. Controls should narrow the available actions and enforce authorization when those actions reach other systems. [11]

A production assistant could, for example, prepare a proposed schedule change for review while the scheduling system preserves authority over the accepted plan. This allows useful preparation without giving an uncertain interpretation direct control over operating commitments.

Automated work also needs sensible limits on volume, duration, and resource use. A fault that repeats a harmless action many times can become disruptive or expensive. Management should know who receives an alert, what can be stopped, and how the organization determines whether an action has already completed.

Protect inputs and calculation integrity

The reliability of a mathematical result depends on the integrity of its inputs and method. A changed source record, unauthorized parameter revision, or substituted calculation package can produce a convincing but unsuitable answer. Security and model governance must therefore meet at the points where those elements can change.

Approved versions should be identifiable, and significant changes should leave a record. Access to modify the model should be distinct from ordinary permission to run it wherever the system supports that separation. Before an updated implementation is accepted, tests should confirm that the intended calculation and operating rules remain intact.

Protection also extends to the supporting software and integrations. Someone must own their maintenance, assess relevant updates, and understand what a change could affect. A secure initial deployment needs continuing attention to remain suitable.

Examine suppliers and service dependencies

An organization's controls extend across services it does not operate itself. When choosing an AI service or integration,

management should examine the relevant information handling, access arrangements, incident support, and dependence on other providers. The review should address the service and configuration actually used.

A statement that business data is excluded from model training answers one question. It does not settle every question about retention, authorized access, connected services, or the company's own use of the output. Each material concern needs an appropriate answer.

The business should also understand how it could continue if a connection became unavailable or a supplier arrangement ended. Important operating knowledge, approved calculations, and accepted records should remain recoverable in a form the organization can use.

Detect problems and restore useful work

Monitoring should make consequential failures visible: unexpected external actions, unexplained changes, unusual access, missing records, or repeated execution failures. Logs only support oversight when someone knows what to examine and how to respond.

An incident procedure should identify who can suspend a workflow, preserve the relevant evidence, involve the appropriate specialists, and coordinate communication. Recovery should establish both that the system is available and that its data and configuration are trustworthy enough to resume the intended work.

Exercises can test this before a real interruption. For an important reporting process, staff might practise preparing the report from the last approved data and method while the normal connection is unavailable. The exercise reveals whether the fallback is usable and whether the responsible people understand it.

NIST's AI Risk Management Framework organizes AI risk management around Govern, Map, Measure, and Manage. This voluntary framework supports attention to trustworthiness throughout AI design, development, use, and evaluation. It provides a broader management context for the security controls and oversight discussed here. [9]

Make security part of business confidence

Effective security can support adoption by clarifying which information people may use and which actions they may delegate. It can also help a business answer reasonable questions from customers about how its information and commitments are protected.

Nessiva's capabilities include cybersecurity and AI security assessments, technology risk, governance, control evaluation, and readiness for relevant assurance work. Our team connects these questions with the analytical model and operating process. An assessment identifies findings and improvements within its scope; any independent certification or formal attestation requires its own appropriate engagement.

This is the contribution of security expertise to the combined approach. It helps the business preserve the value created by faster execution and better analysis, while remaining prepared to detect problems, intervene, and recover.

CHAPTER 16

The Risk and AI Officer function

AI applications and quantitative models need continuing attention if they are to remain useful. A business must decide what to develop, how to validate it, which operating problems deserve priority, and who will keep the work connected to management decisions.

A Risk and AI Officer function can coordinate this responsibility. It brings together the development of the business model, selected AI applications, and the rules under which both are used. Its purpose is to improve business performance while addressing the possibility of unwanted loss.

Why risk belongs in the title

Risk is relevant whenever the business can lose something it values. Lost money is one example. So are avoidable rework, specialist time consumed by preventable problems, unreliable delivery, weakened customer trust, and reduced ability to handle future demand.

A quantitative view helps management examine how these effects interact. An AI application may reduce one loss while introducing another if its output is poorly controlled. The role therefore considers both the intended improvement and the conditions required for that improvement to remain useful.

A coordinated programme

The function can work on several connected priorities at the same time. It may deepen a selected part of the model, develop an AI pilot for an information-heavy task, and establish the controls needed before that pilot reaches wider use.

Some activities can proceed in parallel. Others have necessary dependencies. A team can prepare policy while mapping the

workflow, but it cannot credibly validate a measure before deciding what the measure means. It can prepare test cases before records arrive, but synthetic tests cannot establish real operating performance.

Coordination makes those distinctions explicit. It also gives management a place to resolve conflicts when several departments want attention from the same limited analytical or technical resources.

A fractional role can be proportionate

A smaller business may need continuing specialist leadership without needing a full-time internal department. A fractional arrangement can provide a defined allocation of professional effort, access to relevant expertise, and an agreed reporting cycle.

The arrangement still requires participation from the business. Managers explain priorities, staff help establish how work is actually performed, and designated people review and approve outputs. A fractional function coordinates specialist work; it does not eliminate operating responsibilities.

Its mandate should specify authority, expected outputs, access arrangements, and dependencies. The role can advise and develop tools without becoming the line manager of the people using them. Management retains authority over business direction and implementation.

Report work in a form leaders can question

A useful reporting cycle distinguishes planned activity, completed deliverables, test results, unresolved dependencies, and next decisions. It should make clear whether a tool is a prototype, a controlled pilot, or an approved operating workflow.

A controlled management assistant can make these records easier to question. Leaders can ask what was delivered, why a priority changed, or what remains unresolved. Answers should refer to approved records and preserve the difference between intention and completion.

Management can then judge whether the evidence justifies the next investment.

CHAPTER 17

Measure value in the complete process

A demonstration can show that AI performs a task. A business evaluation must establish whether the resulting workflow produces a worthwhile improvement after preparation, review, correction, and continuing support are included.

The baseline should describe how the work is currently performed and the requirements the result must meet. Time matters, but so do completeness, accuracy, reliability, and the consequences of a missed issue. The evaluation should preserve the same business purpose on both sides of the comparison.

A simple example

The following numbers are hypothetical and describe a comparison method, not a client result.

Suppose a recurring document task occurs 200 times each month. The current process takes 30 minutes of staff time per case, including its existing checks. That requires 100 staff hours per month.

In a proposed workflow, preparation takes 10 minutes per case and review and correction take another 5 minutes. The complete task therefore takes 15 minutes, or 50 hours per month. The potential release is 50 hours, provided the tested quality requirements are met and the cases are representative.

Monthly comparison	Current process	Proposed workflow
Number of cases	200	200
Total staff time per case	30 minutes	15 minutes
Total staff time	100 hours	50 hours
Potential staff time released	Baseline	50 hours

The proposed total includes human review and correction. It does not include one-time development or training, or continuing maintenance and service costs. Those must be assessed separately before deciding whether to invest.

The next executive question is what those released hours allow the business to do. If they reduce overtime, there may be an identifiable cost benefit. If they allow a constrained team to clear valuable work, the benefit may be greater capacity or faster delivery. If the work cannot be reassigned usefully, the immediate commercial benefit may be smaller.

Test quality where failure matters

A workflow can be faster on average and still unsuitable if it misses important issues. Evaluation should consider errors by consequence, including cases where the system reports no problem but a relevant problem exists.

The test set should reflect the work the system will actually encounter. Easy cases alone are insufficient. The team should also examine missing information, inconsistent revisions, unusual formats, and other conditions likely to challenge the workflow.

An experienced reviewer can compare the system's output with an independent assessment. Differences should be investigated and recorded. The objective is to understand when the workflow helps, when it needs revision, and where its limits should be set.

Include the full cost of operating

The cost of integration can include specialist design, staff participation, data preparation, configuration, software access, technical connections, training, testing, review effort, and maintenance. Some costs occur once; others continue.

The proposed benefit should also avoid double counting. Time released from a process and additional output produced using that time are related. Treating both as independent financial gains can exaggerate the result.

Similarly, a faster intermediate step does not prove that customer delivery becomes faster. Another constraint may absorb the gain. The business should measure the final outcome that justified the initiative as well as the local task improvement.

Include the value of effective governance

Governance can improve performance by reducing uncertainty about who may decide and what evidence they need. A well-designed approval can avoid repeated clarification; a maintained source can reduce conflicting answers; a clear exception route can help staff resolve unusual work sooner.

These benefits should be examined alongside the effort of operating the controls. Unnecessary approvals can create delays, and unread logs add little value. The business should assess whether the chosen controls address material risks and support useful work. The objective is accountable execution with a proportionate burden.

Make a decision from the evidence

A pilot should end with a practical recommendation. Continue if the benefit is worthwhile and the required controls are effective. Revise if a specific, repairable issue prevents success. Stop if the application is unsuitable or the expected benefit does not justify the cost.

Stopping a weak application is useful management work. It preserves attention and resources for a stronger opportunity. The value of a pilot includes learning where the technology belongs.

After launch, the same logic continues. Compare actual performance with the expected result, investigate significant changes, and review whether the workflow still serves the business priority. A useful tool earns continuing use through evidence.

CHAPTER 18

Begin with a practical sequence

The first implementation should answer a real business question within a manageable scope, with its responsibilities and safeguards established alongside the analytical work. It should also create knowledge that makes the next decision easier. The following sequence describes the executive logic of the work without prescribing a fixed timetable for every business.

Establish the purpose and sponsor

Begin with the outcome leadership wants to improve and the operating problem believed to affect it. Name the person who can set priorities and resolve access or responsibility questions. Agree what a useful first result would look like.

This need not begin with a full enterprise model. A narrowly defined advisory question or document workflow may be sufficient. If the problem concerns competing priorities across the company, a Type A foundation may be the more useful starting point.

Understand the current work

Meet the responsible manager and the people who perform the task. Examine a small, representative set of records. Establish what the work is for, what determines an acceptable result, where it waits, and where decisions are made.

This stage tests assumptions. A visible delay may result from missing information rather than slow document preparation. An AI tool that drafts faster will not resolve the underlying problem unless the workflow also addresses that missing information.

Define the evidence and boundaries

Agree on the baseline and the proposed measures. Identify the authoritative documents and the limitations of the available data.

Specify what the first model or workflow will cover, what decisions it may support, and which actions require approval.

The scope should be understandable to the people who will use and review it. A phrase such as document review is too broad if different readers assume it includes different checks. State what is being compared and what the output will contain.

Develop analysis and controls together

Build the first useful analytical structure or workflow while resolving data access, review, and change responsibilities. Record its owner and permitted actions, test the relevant controls, and prepare clear instructions and a way for staff to challenge the result. A brief governance arrangement that people can apply is a useful starting point.

Where quantitative development and an operational AI application support the same priority, they can progress together. The model can clarify where the business effect should appear, while the workflow can produce evidence about the practical improvement.

Test before wider use

Use representative cases and assess the complete process. Record the time required, the quality of the output, the corrections needed, and any limitations. Include conditions where the system should ask for help or stop.

When real records are unavailable, constructed or de-identified cases can support early development. Their results must remain clearly distinguished from validation in the actual operation. Management should know which uncertainty has been resolved and which still depends on access to real work.

Decide and extend

Review the evidence with the responsible people. Agree whether to continue, revise, or stop. For continuing work, identify the owner, operating instructions, training, monitoring, and fallback arrangements.

Extend the approach where the evidence supports it. A successful workflow in one document type may require new testing for another. A Type A model may identify a useful Type A-

B investigation. A series of comparable observations may justify a later statistical extension.

Make staff participation practical

Employees need to understand the intended benefit and how their knowledge influences the design. They should be able to identify misleading measures, missing cases, and rules that do not reflect the work.

Training should use the tasks people actually perform. It should show the expected input, a useful output, an example of an error, and the route for getting help. A general presentation about AI does not replace this operating knowledge.

Management should also decide how released time will be used. Giving staff a tool without adjusting expectations or removing duplicate work can leave the workload unchanged. The process needs to change enough for the benefit to reach the people and outcomes it is intended to support.

The sequence is disciplined because each stage supplies something the next stage needs. It is adaptable because the appropriate depth depends on the business problem and the evidence available.

CHAPTER 19

Work with Nessiva

Nessiva is a Toronto-based corporation helping businesses connect mathematical modelling, AI, risk management and governance to the decisions and operating improvements that matter. Our team brings analytical and assurance capabilities together so a business can develop useful tools with clear responsibilities around them.

This approach draws on established professional disciplines and years of experience in business, mathematics, risk, technology, and controls. AI strengthens how we analyse, build, test, and communicate. The business objective determines how that capability is applied.

The capabilities we bring

Our quantitative work centres on *Virgilio* and *Dante*. *Virgilio*, built on the Nessiva Utility Vector Framework, structures organizational evidence and measures business health through performance, capacity, risk and valued outcomes. *Dante* adds statistical analysis, including correlations, regressions, time-series methods and appropriate scenario or forecasting tools.

Our AI work includes identifying worthwhile applications, preparing company-specific kernels and reusable skills, configuring controlled workflows and evaluating their usefulness. We specialize in ChatGPT for Business. Autonomous, deterministic and judgment-embedded kernels support management analysis and practical work with documents, reporting and operating knowledge.

Our risk, policy, and governance work addresses the rules and responsibilities needed to use these capabilities consistently. It includes model oversight, change control, management education, feedback mechanisms, and the examination of business and technology risks.

Our cybersecurity and assurance work addresses information protection, access, technology controls, operational resilience, and readiness for relevant assurance requirements. This includes

security assessments aligned with established frameworks and practical recommendations for management. The purpose is to connect the intended use of AI with the controls it needs.

The value of an integrated team

The benefit of bringing these disciplines together is practical. A mathematical design can be examined for its governance implications while the AI workflow is being developed. A security requirement can be incorporated before an information connection becomes difficult to change. The business receives a coordinated explanation of what the system does and how it should be managed.

Depending on the engagement, tangible outputs can include an analytical model, a tested workflow, a record of responsibilities and permissions, operating policies, assessment findings, and a plan for continuing review. Scope is agreed around the decision and the level of support the business needs.

A structure that grows more useful over time

The mathematical skeleton gives these capabilities a common purpose. It connects the business's chosen outcomes to operating evidence and supplies continuity as AI models and applications develop. A new assistant, integration, or analytical method can then be assessed by the contribution it makes to that structure.

With suitable records, repeated use can support a history of performance, decisions, and results. We help businesses develop the definitions, kernels, calculation methods, and governance needed for that history to remain interpretable. Later analysis can investigate trajectories, delayed effects, and forecasts where the evidence warrants them.

This is the opportunity we want business leaders to recognize: practical improvements in everyday work can form part of an enduring management capability. The first useful application can also prepare the organization for a more connected and informed way of operating.

Three levels of engagement

Level 1 — One-day Ignition. A seminar and hands-on workshop at the workplace helps staff begin the transition to practical AI use.

Specializing in ChatGPT for Business, we combine demonstrations, guided practice with relevant tasks, information-handling guidance and checking of outputs. The team identifies promising applications and practical next steps. Short reviews, focused advisory and brief analyses are also available within our short-term engagements.

Level 2 — Implementation. A scoped project develops and integrates *Virgilio*, *Dante* and selected AI workflows. This may begin with the Type A foundation, deepen through Type A-B work, and add analytics as evidence becomes available. Deliverables can include models, data connections, kernels, tested workflows and operating controls. The scope states the business purpose, required evidence, responsibilities and expected output.

Level 3 — Continuing support. An ongoing or fractional Risk and AI Officer engagement provides continuity through recurring management analysis, model development, AI integration, staff education and review. Governance, cybersecurity and clear decision authority accompany this work, with the mandate and reporting rhythm agreed for the business.

A business can enter at the level appropriate to its needs. A one-day engagement is complete in itself and does not require a longer commitment. Implementation and continuing support are scoped separately. Across all three levels, Nessiva combines practical AI adoption with quantitative expertise, governance and cybersecurity.

Prepare for a useful first conversation

You do not need a complete AI strategy to speak with us. Bring a short description of the business and one consequential problem. Explain where the difficulty appears, who is affected, and what a better result would mean.

If possible, identify the kinds of records already available and the person who understands the work. Confidential documents do not need to be included in an initial enquiry. Appropriate information arrangements can be established as the discussion becomes more specific.

Useful starting questions include whether a bottleneck deserves attention, whether a current measure supports a decision, where AI could reduce recurring effort, or what controls are needed before adopting a proposed workflow.

Our role is to turn that question into a clear assessment and an appropriate next step. For some businesses, the next step will be a focused improvement. For others, it will be a more connected view of the organization and a continuing programme of development.

Visit **nessiva.com** and use the contact form to request an initial consultation. Describe what the business needs to improve, and we can begin with the decision that matters.

About the authors

Maksim Sokolov

Maksim Sokolov is a co-founder and Managing Director of Nessiva. He combines mathematical modelling, quantitative risk analysis, statistics, and applied AI to help business leaders understand performance, evaluate trade-offs, and make decisions about resources and operational improvement.

His business and consulting experience dates back to 2007. He has extensive experience providing quantitative consulting and risk management services to leading retail and production companies. His advisory work focuses on business performance, quantification, capacity analysis, and decision models. It also includes the assessment of business and financial risks, dependencies, and trade-offs, together with the integration of AI into analysis, decision support, and operational workflows.

Maksim holds a PhD in Mathematics and master's degrees in Mathematics and Mathematics Education. He also holds the Professional Risk Manager (PRM) designation from the Professional Risk Managers' International Association (PRMIA).

He is a Professor of Business and Financial Mathematics and Statistics at Seneca Polytechnic's School of Business. His work at Toronto Metropolitan University included financial risk, financial derivatives, and business and financial mathematics. His earlier academic appointments include Cardiff University, the International Centre for Theoretical Physics, the National University of Uzbekistan, and Moscow State University.

Maksim is an author of research articles and books in mathematics, business mathematics, and financial mathematics. At Nessiva, he connects this academic background with the practical demands of business decision-making.

Sanjar Alimov

Sanjar Alimov is a co-founder and Managing Director of Nessiva. His expertise connects cybersecurity and AI security with technology risk, audit, governance, and operational resilience. He helps businesses protect information, oversee AI use, and maintain human accountability.

He brings more than 20 years of experience, including leadership roles in banking and consulting work in professional services. His work spans technology and cybersecurity assurance, AI and model-risk oversight, and the assessment of governance and controls. He translates complex business and technology issues into practical recommendations for senior management and boards.

Sanjar holds a Master of Accounting, an MSc in Economics, and a BSc in Economics.

His professional certifications include Certified Information Systems Security Professional (CISSP) and Certified Cloud Security Professional (CCSP) through ISC2; Certified in the Governance of Enterprise IT (CGEIT) and Certified Information Systems Auditor (CISA) through ISACA; Certified Internal Auditor (CIA) and Certification in Risk Management Assurance (CRMA) through the Institute of Internal Auditors; and the GIAC Critical Controls Certification (GCCC).

His experience includes IT audit planning, controls assessments, data analytics, executive reporting, and multidisciplinary team leadership. His background also covers identity and access management and cloud and hybrid technology environments.

At Nessiva, his advisory focus includes technology oversight, control effectiveness, compliance, business continuity, and disaster recovery. His security and assurance work covers NIST- and ISO-aligned control reviews, audit readiness, customer assurance, data protection, access controls, and human oversight of AI.

Sources and notes

Scope of the sources

The business analysis and proposed operating designs are the authors' interpretation. Public sources support the specific findings, product descriptions, and established frameworks cited in the text. Implementation examples are anonymous, and illustrative applications do not report client outcomes.

Product descriptions were checked against official OpenAI documentation on 11 September 2026. Plans, models, interfaces, integrations, and limits can change. An administration guide does not establish that every control is available under every subscription.

References

1. Generative AI at Work. Erik Brynjolfsson, Danielle Li, and Lindsey Raymond. Working paper, revised November 2024, arXiv version 2. The study concerns a particular customer-support deployment. Its results are included as evidence of possible gains and variation across users, not as a projected result for Nessiva clients.

https://arxiv.org/abs/2304.11771v2

2. Models. OpenAI. Current descriptions of GPT-6 Astra and the GPT-5.6 family, their intended uses, and availability qualifications.

https://learn.chatgpt.com/docs/models

3. Use ChatGPT. OpenAI. ChatGPT Work, projects, files, tools, and reviewable outputs.

https://learn.chatgpt.com/docs/use-chatgpt

4. Introduction to Time Series Analysis. NIST/SEMATECH e-Handbook of Statistical Methods. A reference for the established statistical study of observations over time.

https://www.itl.nist.gov/div898/handbook/pmc/section4/pmc4.htm

5. Pricing. OpenAI. Plan descriptions, Business workspace features, and the distinction between subscription and usage arrangements. This book does not quote prices.

https://learn.chatgpt.com/docs/pricing

6. Plugins; Plugins and skills. OpenAI. Supported integrations and practical examples involving email, calendars, and reusable procedures. Each connection's available actions and permissions require verification.

https://learn.chatgpt.com/docs/plugins

https://learn.chatgpt.com/training/walkthroughs/plugins-and-skills

7. GPTs and Sharing. OpenAI. GPT ownership, sharing, connected apps, custom actions, and their administration boundaries. The presence of an enterprise administration guide does not establish that every control is available on every Business plan.

https://learn.chatgpt.com/docs/enterprise/gpts-and-sharing

8. Scheduled tasks. OpenAI. Recurring tasks, supported event triggers, execution conditions, and availability limitations.

https://learn.chatgpt.com/docs/automations

9. AI Risk Management Framework. National Institute of Standards and Technology. The framework is a voluntary public reference for AI risk management. The accompanying Generative Artificial Intelligence Profile addresses risks associated with generative AI. These resources inform general risk-management discussion; no endorsement of Nessiva is implied.

https://www.nist.gov/itl/ai-risk-management-framework

https://doi.org/10.6028/NIST.AI.600-1

10. Prompt Injection. OWASP Gen AI Security Project, LLM01:2025. A reference for risks from instructions in user input and external content.

https://genai.owasp.org/llmrisk/llm01-prompt-injection/

11. Excessive Agency. OWASP Gen AI Security Project, LLM06:2025. A reference for limiting functions, permissions, and autonomy in connected AI systems.

https://genai.owasp.org/llmrisk/llm062025-excessive-agency/

The basis of the book

The presentation of Nessiva's capabilities and framework is grounded in the firm's published descriptions of its services, Utility Vector Framework, progressive model development, and focused advisory. The author profiles follow the firm's published biographies. Website material was reviewed in September 2026.

https://nessiva.com

https://nessiva.com/model.html

https://nessiva.com/advisory.html

The anonymous manufacturing implementation account is grounded in completed model documentation and related engagement materials. Continuing development is described as such. Confidential records are not reproduced or linked, and no confidential client identity is included.