Loop XXI White Paper

Reviewed 2026-11-05

Abundance makes verification scarce

AI makes it easier to produce a draft, an analysis or a piece of software. A business still has to decide whether the result is correct, useful and worth the resources it took to produce. A faster first attempt can leave more work for the person who checks it. A convincing report can conceal a weak assumption. A new tool can introduce another dependency before it removes an old task.

When something becomes cheap to produce, its value shifts to what stays scarce and still useful. For a business adopting AI, the scarce input is verification: the attention it takes to check a result, and the evidence that shows whether a change helped. Drafts are now plentiful. Results that someone has checked are not.

Loop XXI's thesis is that verification deserves the discipline given to any other scarce resource: define the work, understand its cost, inspect the result and keep responsibility clear. The useful unit of progress is a task completed to an agreed standard. Counting generated answers or installed tools tells us little about that standard.

The same principle applies to money. A claim is easier to rely on when anyone can check it, and the supply of a money is a claim like any other. Bitcoin is the clearest working case: its issuance follows a published rule that anyone can verify. That is why Bitcoin sits at the center of how Loop XXI treats its own long-term capital, alongside cash for obligations that fall due sooner.

This paper sets out the reasoning behind that approach. It connects four concerns that are often reviewed separately: how work gets done, how its effects become visible, how the underlying systems are operated and what money the results are kept in. AI adoption, charts and reporting, balance sheet management and Bitcoin meet at the same question: what should change, and what evidence would show that the change helped?

Start with a task someone can check

A useful AI evaluation begins with a particular task. Drafting a follow-up after a sales call, answering an order question and matching a receipt to a transaction each impose different demands on the person checking the result. Grouping them under a single promise of productivity hides the differences that determine whether a tool is useful.

Loop XXI's published work on time savings starts with the time required to do a task by hand, then accounts for checking and redoing the tool's work. The central relationship is simple: a tool has room to save time when checking its output is substantially quicker than producing that output from scratch. When checking requires repeating most of the work, even a strong tool may offer a small gain.

Published studies show both sides of that relationship. Among 5,172 customer support agents, an AI assistant raised issues resolved per hour by 15 percent on average, and newer agents gained the most. Among 453 professionals doing writing tasks, those using ChatGPT took 40 percent less time and their work was graded 18 percent higher. Among 758 knowledge workers, those using AI completed 12.2 percent more tasks and worked 25.1 percent faster on tasks the AI handled well, but were less likely to reach a correct answer on a task just outside its ability. The gains are real and uneven, so the person checking has to know where the tool fails.

The practical test is to compare representative work with and without the tool. Record the time spent, the results accepted without a rewrite and the errors discovered during review. Include ordinary cases and awkward ones. A demonstration selected because it works well cannot establish whether the tool will handle the work that arrives next week.

Small trials can identify a clear difference. Close results need more evidence, particularly when the cost of an error is high. The decision should reflect that uncertainty. A promising experiment justifies further testing; it does not establish that every related task should be automated.

Count the cost of getting it right

The cost of an AI attempt is only one part of the cost of useful work. A person spends time reviewing successes and failures alike. Unusable attempts require another attempt or a manual correction. Setup, training and moving away from a tool also consume resources, even when they do not appear on the same invoice.

Loop XXI's cost-per-successful-task article uses a deliberately simple comparison: add the cost of an attempt to the cost of checking it, then divide by the share of attempts that succeed. That makes reliability visible in the comparison. A tool that costs more per attempt can be competitive when it produces more usable results and reduces repeated checking.

The calculation has a boundary. It assumes failures are caught and cannot account by itself for the consequences of an error reaching a customer or the books. Those risks need separate consideration. A customer-facing commitment and an internal first draft should not be judged by the same tolerance for mistakes.

Measured results also have a date. Changes to a model, a workflow or the person reviewing it can change the economics. A sensible evaluation leaves a record of what was tested and why it passed. That record gives the business a basis for comparison when the tool or the work changes.

Make reports answer a decision

Reporting should make a decision easier to examine. A chart earns its place by showing a relationship that matters: how long cash may last under stated assumptions, where spending is concentrated or whether a process is improving. Adding a chart without a

decision in mind can make a report longer while leaving the underlying question unanswered.

The reader needs to distinguish records from assumptions and calculations from forecasts. A cash projection depends on the money expected to come in and go out. An apparent improvement in task completion may depend on which tasks were included. A useful report makes those dependencies visible so another person can question the result.

This is why formulas and source records matter alongside presentation. When a number changes, the reader should be able to understand whether the business changed, the inputs changed or the calculation changed. A clear visual treatment helps, but it cannot repair an undefined measure.

Loop XXI's published articles include worked examples with stated assumptions and formulas, and link to research records. Those examples explain a method. They are not evidence that a particular business will achieve the same result. Their purpose is to give readers something they can inspect and replace with their own measurements.

Connect operating choices to cash

AI and infrastructure decisions eventually meet the balance sheet. Software can create recurring costs and maintenance work. An implementation can consume time before it produces a benefit. A business may have an attractive operational idea and still need to understand whether its available cash and working capital can support it.

Looking at cash runway, current assets and liabilities, and the timing of collections and payments gives those choices context. The questions are concrete: what obligations fall due, what resources are available and which assumptions could change the picture? A single ratio cannot answer all of them. Each describes part of the business and depends on the quality of the underlying records.

A financial tool should let a reader inspect how its output follows from the inputs. Showing the formula helps reveal what a measure includes and what it leaves out. Keeping a projection beside its assumptions lets the reader consider a different scenario without confusing it with an observed result. The purpose is to make the financial picture easier to examine.

A financial review also needs a clear scope. Understanding obligations, costs and liquidity is different from making an investment decision. Keeping that distinction visible helps the reader use an analysis for the question it can answer, without extending its conclusions into questions it has not examined.

A supply rule anyone can check

A balance sheet counts dollars. It does not show what those dollars will buy, or how fast new dollars are created. From January 2007 to September 2026, the U.S. money supply measured by M2 more than tripled. Over the same period, consumer prices rose 65 percent and average hourly pay rose 84 percent. The longer record shows the same

pattern: from January 1964 to September 2026, M2 grew 59-fold, consumer prices 11-fold and the hourly pay of production and nonsupervisory workers 13-fold.

M2 growth is not the same as lost purchasing power. Consumer prices measure that loss, and they rose far less than M2. Under the equation of exchange, money growth beyond growth in the overall price level, a broader measure than consumer prices, shows up as growth in real output, slower circulation of each dollar, or both. What M2 does show is that the supply of dollars changes at a pace no published rule sets. For six decades, M2 grew between roughly zero and 14 percent a year. It grew 26.8 percent in the year to February 2021, then stayed below its level of a year earlier for fifteen months, from December 2022 to February 2024.

Bitcoin's supply follows a published schedule. Each block paid 50 bitcoin at launch, and that amount halves every 210,000 blocks. Under the current consensus rules, total supply approaches 21 million and never exceeds it. Annual supply growth fell from 14.9 percent in 2013 to 0.8 percent in 2025. The schedule slightly overstates the coins that can be spent, because a few blocks claimed less than the full amount and some coins are lost. Anyone who runs the software can verify the rule and the current supply. Verification is open to everyone, but it is not automatic: a holder who does not run the software relies on someone who does.

A predictable supply is not a predictable price. In the price series Loop XXI's research uses, Bitcoin fell 83 percent from its 2013 peak, 79 percent from its 2017 peak, 76 percent from its 2021 peak and 53 percent from its October 2025 peak to June 2026. From November 2021 to September 2026, its price moved more than 10 percent in 30 of 59 months. Consumer prices did not move more than 1.3 percent in any month of that period. Scarcity limits supply. It guarantees neither a return nor a stable short-run value.

Cash and Bitcoin therefore do different jobs. Cash suits obligations that fall due within months, because its short-run purchasing power is stable. Loop XXI's case for holding Bitcoin rests on a long horizon and a supply no one can expand at will. Loop XXI treats Bitcoin as its long-term treasury asset. Cash covers operating costs and near-term obligations first, so the company does not depend on selling Bitcoin to pay a bill. This describes Loop XXI's own capital.

Infrastructure should make its behavior inspectable

A business depends on more than the quality of a model's answers. It also depends on whether information can be found, whether requests can be understood and whether the systems involved expose enough evidence to check what happened. Infrastructure becomes easier to operate when its inputs, outputs and limits are clear.

The useful pattern is a narrow question followed by a result that can be checked. Can a website be found and read? What does a payment request ask the payer to authorize? How does a financial measure follow from the supplied figures? Making the question explicit limits the temptation to mistake a convenient interface for a complete assessment.

A working tool demonstrates a particular capability. It cannot establish that every business needs the same system. A review may point to a simpler report, a removed subscription or a better-defined checking step. The appropriate change depends on the problem. Adding software creates an obligation to understand and maintain it.

Control also has practical boundaries. Reading information, explaining a result and acting on it require different authority. A well-defined system makes those boundaries understandable to the person using it. This is especially important when a workflow moves from producing a draft to changing records, sending a message or committing resources.

Current work

Loop XXI. Business consulting and infrastructure in the age of AI.

AI adoption, Charts and reporting, Balance sheet management.

We research your infrastructure and finances and send a written list of the problems we find, ranked by impact.

We quote fixes for the problems we can solve. You choose which ones we do. The review and any fixes are quoted and billed separately.

Products

Request a review

Scope and responsibility

We do not hold customer funds or wallet keys. We do not ask for a seed phrase or a private key.

We do not recommend investments or manage anyone's money.

Review first, choose the work second

A review should help a customer decide what work is worth undertaking. Diagnosis and implementation answer different questions. The first asks what is wrong and how much it matters; the second asks what a proposed change would require. Keeping those decisions separate makes room to compare an intervention with other uses of the same time and resources.

That separation matters because finding a problem does not settle how it should be solved. A ranking gives the customer a basis for choosing where to spend attention and resources. The right response may be a process change, a focused implementation or a decision to leave a system alone until better evidence is available.

The framework in this paper is a way to examine those choices. Start with the actual work. Count checking and failure alongside production. Make the financial assumptions

readable. Prefer systems whose behavior can be inspected. Define what evidence would justify continuing, changing course or stopping.

AI adoption becomes more useful when it is connected to the ordinary disciplines of running a business. The question is whether people can complete worthwhile work reliably, understand what it costs and make the next decision with better information. Loop XXI applies the same test of checkable rules to its own capital. That is the standard proposed here for Loop XXI's consulting, research and infrastructure.

Sources

Loop XXI consulting

Loop XXI: Cost per successful task

Loop XXI: Where AI saves time in a small business, and where it does not

Loop XXI products

Loop XXI About

Brynjolfsson, Li and Raymond: Generative AI at Work, QJE 2025

Noy and Zhang: Productivity effects of generative AI, Science 2023

Dell'Acqua and others: Navigating the Jagged Technological Frontier, HBS 2023

Loop XXI: The balance is not the buying power

Loop XXI: Bitcoin and liquidity, what the data can tell us

Loop XXI Research, Lesson 1: The Problem

Loop XXI: The gap between pay and money since 1964

FRED: M2 (M2SL), Board of Governors of the Federal Reserve System

Bitcoin.org: frequently asked questions