TIME

The Standard for Measuring Value in the Age of Autonomous Intelligence

---

Ash Koosha

Original thesis: June 2025
This edition: February 2026

ABSTRACT

Every economic system in history has been, at its root, a system for allocating time. Gold represented stored labour time. Currency represented a transferable claim on another person's hours. Equity represented a bet on future time deployed productively. In each case, money was the abstraction layer and time was the underlying reality.

The rise of autonomous artificial intelligence—systems that perform cognitive and operational work without human intervention—has made this underlying reality impossible to ignore. For the first time in economic history, value is being created by non-human actors operating in time but not consuming human time. The output of these systems is not primarily products or revenue. It is human time, liberated.

Yet no economic framework measures this. GDP counts production. Revenue counts transactions. Productivity metrics count output per input. None of them count the most important thing the intelligence revolution produces: hours returned to human beings.

This paper introduces TIME: a universal standard for measuring, scoring, and transacting the economic value of time in a world run by autonomous systems. We establish the mathematical foundations of the Time Delta Metric (TDM), Quality-Adjusted Time Units (QATUs), the Temporal Signal-to-Noise Ratio (TSNR), and the Consumption-Productivity Index (CPI-T). We

propose that time is not a derivative of value—time is value itself—and that the correct response is to build the measurement infrastructure that makes it visible, auditable, and actionable.

Part I: Why Time

1.1 The Hidden Primitive

Consider what happens when a financial analyst values a company. They examine revenue, margins, growth rate, market share, competitive position, and future cash flows. They discount those cash flows to present value and arrive at a number denominated in dollars. At no point in this process does anyone ask: how much human time does this company liberate or consume?

This omission is not an oversight. It is a structural blindness inherited from an era when all economic value was produced by human effort. In that world, measuring money was sufficient because money was a reliable proxy for time—specifically, for the time humans spent producing things. A dollar earned roughly corresponded to a unit of human time deployed.

That correspondence has broken. Autonomous systems now produce economic output—drafting contracts, diagnosing conditions, routing logistics, generating creative work, managing communications—without consuming human time. The money still flows, but the link between money and human time has been severed. We are measuring the shadow on the wall while the fire has moved.

Time was always the base layer. Money was always the abstraction. The abstraction was useful when the two moved together. They no longer do.

1.2 A Brief History of Time as Capital

The relationship between time and economic value is not new. It is the oldest relationship in economics, merely the most neglected.

In 1776, Adam Smith observed that the real price of everything is the toil and trouble of acquiring it. Labour, in Smith's framework, was the "real measure of the exchangeable value of all commodities." But labour is nothing more than human time applied to a task. Smith was measuring time. He lacked the vocabulary to say so.

David Ricardo formalised this into the labour theory of value: the value of a good is determined by the total labour time required for its production. Karl Marx extended this into surplus value—the difference between the time a worker labours and the time for which they are compensated. The entire Marxist critique of capitalism is, at its root, a dispute about the fair allocation of time.

The marginalist revolution of the late nineteenth century—Jevons, Menger, Walras—shifted the frame from production time to subjective utility. Value, they argued, derives from the satisfaction a good provides, not the time required to make it. This was an advance. But it obscured something: utility itself is experienced in time. An hour of deep focus is worth more than an hour of distracted consumption. The marginalists measured the intensity of experience but lost the temporal dimension in which experience occurs.

In 1930, John Maynard Keynes made a prediction in his essay Economic Possibilities for our Grandchildren. He foresaw that technological progress would reduce the necessary working week to fifteen hours within a century. By 2030, he believed, the central economic problem would not be scarcity but the management of abundance—specifically, the abundance of free time.

Keynes was right about the technology. He was wrong about the distribution. The productivity gains of the twentieth century were absorbed by consumption expansion, not time liberation. We built faster machines and used them to produce more things, not to work less. The economic system captured the time dividend and converted it back into production.

Artificial intelligence changes this dynamic for a structural reason: AI does not merely make human workers faster. It replaces the need for human time entirely in an expanding set of domains. The time dividend can no longer be fully recaptured by production because the production itself requires no human time. The freed time must go somewhere. The question is whether we will measure it, value it, and distribute it—or whether it will remain invisible, as it has for centuries.

1.3 Existing Time Valuations and Their Limits

The idea that time has economic value is not new to this paper. Several systems have attempted to measure, price, or transact time directly. Each succeeded partially and failed instructively. Understanding where they stopped is necessary for understanding where TIME begins.

Time Banking

In 1973, Teruko Mizushima developed the first time banking system in Japan. Edgar Cahn later popularised the concept in the United States through Time Dollars. The principle is direct: one hour of service equals one credit, regardless of what service is performed. A brain surgeon's hour is worth exactly the same as a gardener's hour. Millions of people across dozens of countries have transacted in time credits—from Ithaca Hours in New York (where one Hour note was pegged at $10, the local average hourly wage) to Fureai Kippu in Japan, a system for elder care where people earn credits by helping the elderly and redeem them decades later for their own care or transfer them to family members.

Time Banking proves something important: people will accept time as a unit of exchange. Fureai Kippu proves something further: time credits can function as a store of value, transferable across generations. These are not theoretical claims. They are operational systems with decades of data.

Where Time Banking stops: it refuses to differentiate quality. One hour equals one hour is a political commitment to equality, not an economic measurement. In practice, an hour of emergency cardiac surgery produces a different magnitude of state change than an hour of lawn maintenance. Time Banking addresses this deliberately—it is a solidarity system, not an economic instrument. TIME departs here. The Utility coefficient, the Distribution coefficient, and the Signal-to-Noise framework exist precisely to measure the qualitative differences that Time Banking intentionally ignores. Time Banking demonstrated that time is transactable. TIME provides the measurement that makes the transaction accurate.

Socially Necessary Labour Time

Karl Marx's Labour Theory of Value proposed that the "true" value of any commodity is determined by the socially necessary labour time required to produce it—not the actual time any individual worker spent, but the average time required under normal conditions with average skill and prevailing technology. In this framework, every price is a time score in disguise. Money is congealed labour time. The market is a system for exchanging frozen human hours.

This is remarkably close to the thesis of this paper. The departure—and it is the critical departure—is directional. Marx measured time consumed in production. TIME measures time liberated from production. Marx looked at the input side: how many human hours went into this commodity. TIME looks at the output side: how many human hours does this system give back.

In a pre-AI economy, these are two perspectives on the same quantity. The time consumed in producing a good roughly corresponds to the time burden imposed on the buyer, who must work some number of hours to earn the money to purchase it. But in a post-labour economy, the correspondence breaks. An autonomous system may consume near-zero human labour time in its operation while liberating thousands of hours for its users. Marx's framework cannot capture this because it is anchored to production inputs that no longer require human time. TIME captures it because it is anchored to liberation outputs—the hours returned. Marx described the economy that was. TIME describes the economy that is forming.

Value of Time in Transport Economics

The most direct precedent for TIME exists in transport economics, where governments already assign precise monetary values to saved time and use those values to justify infrastructure spending worth billions. The UK Department for Transport publishes official Value of Time (VoT) figures. The US Department of Transportation maintains equivalent calculations. Every major infrastructure project in the developed world—every highway, railway, bridge, and transit system—includes a cost-benefit analysis in which saved minutes are converted to economic value.

The methodology is direct: economists calculate a VoT, typically set as a percentage of a region's average hourly wage. If a new bridge saves 10,000 commuters 10 minutes per day, the system scores that as a multi-million-dollar annual economic gain, treating those saved minutes as measurable capital. This is TIME in embryonic form—already accepted by governments, already embedded in public resource allocation, already used to move billions in infrastructure investment.

Where VoT stops: it treats all saved time as economically equivalent. Ten minutes saved is X dollars, regardless of what happens with those ten minutes, who saves them, or whether they are used productively or dissipated. VoT has no utility weighting—ten minutes saved from a commute scores the same whether the commuter uses it for focused work or stares at a wall. It has no signal-noise distinction—it cannot tell whether the saved time produces state change or is consumed by noise. It has no distribution coefficient—a bridge that saves wealthy suburban commuters ten minutes scores identically to one that saves low-income workers the same amount.

TIME corrects each of these omissions. It is VoT made complete: time valued not just by quantity but by quality, not just by aggregate but by distribution, not just by the saving but by what the saving produces. Governments already price time. They do it without accounting for utility, signal quality, or equity. The TIME standard provides the full framework.

1.4 The Post-Labour Inversion

We define the Post-Labour Inversion as the economic phase transition in which the primary output of technological systems shifts from goods and services to liberated human time.

In a pre-inversion economy, technology increases the output per unit of human time. A tractor lets one farmer do the work of fifty. A spreadsheet lets one analyst process the data of a department. The human is still in the loop. Time is still the input.

In a post-inversion economy, technology removes the human from the loop entirely. An autonomous agent handles a professional's client communications, scheduling, invoicing, and lead generation without human intervention. The output is the same. The human time input is zero. The "product" of this system is not the communications or the invoices—those are outputs. The product is the human hours that are no longer required.

This inversion demands a new measurement framework. Measuring the outputs (emails sent, invoices processed) tells you what the machine did. Measuring the time liberated tells you what it did for humanity. These are different quantities, and only the second one matters for understanding the true economic impact of autonomous intelligence.

Part II: The Mathematics of Time

2.1 Core Definitions

We establish the following foundational quantities:

Human Time (H). The number of hours a task, process, or workflow would require if performed entirely by a competent human professional, under standard working conditions, without AI assistance. This is the counterfactual baseline. H is measured in hours and must be estimated using industry benchmarks, historical data, task analysis, or expert assessment.

Autonomous Time (A). The wall-clock time consumed by an AI system performing the equivalent task. A includes inference time, orchestration overhead, queue waiting, and any human review or correction time that remains in the loop. A is measured in hours.

Computational Time Cost (C). The total computational resources consumed by the AI system, expressed in time units. This includes GPU hours for inference, KV cache occupancy duration, orchestrator processing time, model loading latency, and any training time amortised across the total expected uses of the model. C is measured in normalised compute-hours.

The relationship between these quantities establishes the fundamental equation of the TIME standard:

$$T = H - A$$

Where T is the raw time liberated—the difference between what a human would have spent and what the autonomous system actually consumed of human attention. This is the simplest expression of time value. Every elaboration that follows is a refinement of this core identity.

2.2 The Time Delta Metric (TDM)

Raw time liberated (T) is necessary but insufficient. An hour saved from a life-threatening medical delay is not equivalent to an hour saved from email formatting. A system that saves one executive a thousand hours is not equivalent to a system that saves a thousand people one hour each, even though the raw total is identical. To capture these distinctions, we introduce the Time Delta Metric:

$$TDM = \Sigma\ (H_i \times U_i \times D_i)$$

Where, for each function or task i:

- H_i = Human hours saved by AI function i
- U_i = Utility coefficient (0 to 1), measuring the qualitative value of the saved time
- D_i = Distribution coefficient (0 to 1), measuring how equitably the time savings are distributed across the user population

The Utility Coefficient (U) captures the qualitative weight of saved time. Not all hours are equal. We propose a sector-weighted baseline:

Domain	Baseline U	Rationale
Emergency medicine and triage	0.95	Time directly correlates with survival outcomes
Education and learning	0.85	Compounding returns on human capability
Legal and regulatory compliance	0.75	High-value professional time with broad impact
Business operations and admin	0.60	Necessary but lower individual impact per hour

Entertainment and media	0.30	Subjective value, limited compounding returns
Addictive engagement loops	0.00–0.05	Near-zero or negative utility

These baselines are adjustable through governance and empirical calibration. The principle is that U must be determined by the nature of the time being returned, not by the revenue generated by the system returning it. A social media platform generating billions in revenue while trapping users in engagement loops produces high monetary value and near-zero time utility.

The Distribution Coefficient (D) captures equity. A system that saves time only for those who can afford premium AI tools scores lower than one that democratises time savings across economic strata. D is calculated as:

$$D_i = 1 - \text{Gini}(\text{time savings distribution across user base})$$

Where Gini is the standard Gini coefficient applied to the distribution of time savings among all users or potential users of the system. A perfectly equitable distribution yields D = 1. A system where all time savings accrue to a single user yields D approaching 0.

2.3 Quality-Adjusted Time Units (QATUs)

Drawing from the QALY (Quality-Adjusted Life Year) framework in health economics—which transformed how healthcare systems allocate resources by weighting survival years by quality of life—we define the QATU:

$$1 \text{ QATU} = 1 \text{ hour of human time liberated at } U = 1.0 \text{ and } D = 1.0$$

A QATU is the gold standard of time value: one hour returned to humanity with maximum utility and perfect distribution. All real-world time savings are expressed as fractions of this ideal.

Example: An AI legal drafting system deployed across 500 law firms saves an estimated 1,000,000 human hours annually. The utility coefficient for legal work is 0.75. The distribution analysis shows the system is accessible primarily to mid-to-large firms, yielding a D of 0.65.

$$\text{QATUs} = 1{,}000{,}000 \times 0.75 \times 0.65 = 487{,}500 \text{ QATUs}$$

This firm produces 487,500 quality-adjusted time units annually. This number is directly comparable to any other system's QATU output, regardless of sector, revenue model, or technology stack.

2.4 The Computational Time Stack

A complete accounting of time value must include the computational time consumed to produce that return. Every AI output carries a time cost that is invisible in current economic analysis.

We define the Computational Time Stack (CTS) as:

$$CTS = T_{train} + T_{infer} + T_{orch} + T_{cache} + T_{latency}$$

Where:

- **T_train** = Amortised training time: total GPU-hours to train the model, divided by expected lifetime uses
- **T_infer** = Inference time: GPU-hours consumed per request or task completion
- **T_orch** = Orchestration time: overhead from agent coordination, tool calls, routing, retry logic
- **T_cache** = Cache occupancy: KV cache and memory utilisation duration during processing
- **T_latency** = Latency overhead: queue waiting, network round-trips, cold-start delays

The Time Leverage Ratio (TLR) expresses the efficiency of an AI system in pure time units:

$$TLR = T / CTS$$

A system with a TLR of 5,000 returns five thousand hours of human time for every compute-hour consumed. This is the most important efficiency metric in the autonomous economy because it is denominated entirely in time, requiring no currency conversion, no market assumptions, and no subjective valuation. It is a physical ratio of time returned to time consumed.

Current frontier models achieve extraordinary TLRs. A single inference call consuming 0.003 compute-hours can draft a document that would take a human 4 hours, yielding a TLR of approximately 1,333 for that interaction. Across millions of interactions, the aggregate TLR represents the raw time-productive power of the system.

2.5 Time Compounding

Time, unlike money, cannot be stored. But it can compound. When an AI system saves a human 4 hours, and that human uses one of those hours to deploy an additional agent that saves 10 other people 4 hours each, the total time liberated is not 4 hours but 44. If those 10 people each deploy their own agents, the cascade continues.

We model this as the Time Compounding Rate (TCR):

$$TCR = T_{generation}(n+1) / T_{generation}(n)$$

Where T_generation(n) is the total time liberated in generation n of a deployment cascade. A TCR greater than 1 indicates positive time compounding—the system creates time faster than it is consumed. This is the temporal equivalent of compound interest, and it is the mechanism by which autonomous agents could produce exponential returns in human time.

No existing economic model accounts for time compounding. Financial compounding is understood. Knowledge compounding is studied in endogenous growth theory. But the compounding of liberated time through cascading autonomy is a new phenomenon that only becomes visible when time itself is the unit of measurement.

2.6 Temporal Entropy

Not all liberated time is structurally equal. An hour returned as a single contiguous block has different value than an hour returned as twelve five-minute fragments across the day. Fragmented time has high entropy—it is difficult to invest in deep work, creativity, relationships, or sustained focus. Contiguous time has low entropy—it can be meaningfully deployed.

We define Temporal Entropy (E) as:

$$E = -\Sigma\ (p_i \times \log_2(p_i))$$

Where p_i is the proportion of total time saved that occurs in contiguous block i. A system that saves 4 hours as a single block has minimal entropy. A system that saves 4 hours as 48 five-minute interruptions has maximum entropy.

The entropy-adjusted QATU (eQATU) incorporates this:

$$\text{eQATU} = \text{QATU} \times (1 - \text{E\_normalised})$$

Where E_normalised scales entropy to the [0,1] range. This penalises systems that technically save time but fragment it into unusable segments—a pattern common in notification-driven tools that save seconds per interaction but impose constant context-switching costs.

2.7 Dynamic Baselines and Temporal Deflation

A measurement standard must remain valid as the systems it measures improve. As GPUs accelerate, as inference approaches real-time latency, and as AI becomes embedded in every professional workflow, the human counterfactual—the H in our equations—faces a stability problem.

Today, H is grounded in observable reality. We know how long it takes a lawyer to draft a contract without AI because lawyers were doing so without AI two years ago. But as autonomous systems saturate professional work, the human-only baseline becomes historical, then hypothetical, then archaeological. Within a decade, asking “how long would a human take without AI” will resemble asking how long it takes to cross the Atlantic by sailing ship—technically calculable, but disconnected from anyone’s lived experience.

The TIME standard addresses this through a rolling relative baseline:

$$\text{H\_relative}(t) = \text{A\_best}(t - 1)$$

The human baseline at time t equals the best autonomous performance at time t minus one recalibration period. QATUs measure the delta between current performance and previous best, not the delta between current performance and a hypothetical AI-free human.

This has three consequences. First, QATUs become measures of marginal time liberation—the improvement over the previous state of the art rather than the improvement over a static historical baseline. This rewards continued advancement rather than rewarding every system equally for outperforming a human with a pen and paper.

Second, QATUs naturally deflate over time. The first AI legal tool that reduced contract drafting from 40 hours to 4 hours generated 36 hours of raw liberation. The next tool that reduces it from 4 hours to 2 hours generates 2 hours. This deflation is correct: it reflects diminishing marginal returns at the frontier of optimisation and prevents QATU inflation as AI proliferates.

Third, to maintain a constant QATU output, a system must keep improving. Standing still for one recalibration cycle means your QATUs drop toward zero because you are no longer advancing beyond the baseline. The standard penalises stagnation and rewards genuine progress—which is how any measurement of technological value should behave.

The recalibration cycle should be annual for sector-level baselines, with the option for quarterly adjustment in fast-moving domains. Each cycle, the previous period's best-in-class performance becomes the new H_relative. This creates a ratchet effect that keeps the standard relevant regardless of how fast the underlying technology accelerates.

On the computational side, as inference costs fall and GPUs become more powerful, the Computational Time Stack (CTS) shrinks toward zero. This causes the Time Leverage Ratio (TLR) to trend toward infinity, which is mathematically correct but diagnostically useless—a metric that approaches infinity for everything differentiates nothing. The TIME Score handles this by design: TLR carries the lowest default weight (0.15), and this weight should decay further as compute commoditises. When computational time becomes effectively free, the remaining three components—QATU, TSNR, and CPI-T—absorb the full weight. The standard adapts to a world of abundant compute by shifting emphasis to what compute produces, not what it costs.

Part III: Signal and Noise in Time

3.1 The Signal-to-Noise Problem

Steve Jobs was described by the people who built products alongside him as having an extraordinary signal-to-noise ratio. He could walk into a room of a hundred options and discard ninety-nine of them in minutes—not because he was dismissive, but because he could see which actions would produce a state change in the world and which ones were motion without movement. He operated almost entirely in the present tense: not what might matter next quarter, but what matters right now, and whether the thing in front of him moved the needle or wasted the room's time.

This quality is not a personality trait. It is an economic force. Every technology, every company, every agent, every system has a signal-to-noise ratio in time. Some produce focused output that changes the state of the world. Others generate enormous activity that leaves everything exactly where it started. The entire AI industry is producing both at unprecedented scale, and there is no metric to tell them apart.

The distinction between signal and noise in time is the most important unmeasured quantity in the modern economy. This section provides the measurement.

3.2 Defining Signal: The State Change Test

Signal is any time expenditure that produces a measurable state change. Something is different after than before. This is the universal test. It applies to a human, an agent, a model, a company, a GPU, and a civilisation. If the state of the world changed in a verifiable way, time produced signal. If it did not, it did not.

State change decomposes into three measurable dimensions:

Delta Output (ΔO). A tangible artefact was produced or a task reached completion. A document was written. A diagnosis was made. A route was optimised. A decision was reached. A transaction was settled. This is the most directly measurable dimension: did something get finished? It is binary at the task level and quantifiable at the aggregate level. ΔO is measured as the time that contributed to producing completed outputs.

Delta Knowledge (ΔK). Information was acquired, synthesised, or transferred that did not exist in the recipient before. A person learned a new skill. A model was fine-tuned on new data. A database was updated with previously absent records. A team reached shared understanding of a problem they previously saw differently. ΔK is measurable through demonstrated capability change, information entropy reduction, or verified knowledge transfer. It is measured as the time that contributed to knowledge state changes.

Delta State (ΔS). The condition of a system, person, or process moved from one state to a different, intended state. A patient's health improved. A codebase was refactored from fragile to robust. A relationship deepened. Infrastructure was upgraded. A mental state shifted from confusion to clarity. ΔS captures transitions that are neither output nor knowledge but are real, verifiable changes in the state of the world. It is measured as the time that contributed to state transitions.

The total signal produced by any system or process is:

$$\text{Signal (S)} = \Delta O + \Delta K + \Delta S$$

Each component is denominated in time units—specifically, the hours that contributed to producing that category of state change. Any time expenditure that produces a non-zero value in at least one dimension is signal. The magnitude of signal is the aggregate state change across all three.

3.3 Defining Noise: The Four Components

Noise is time expenditure where $\Delta O = 0$, $\Delta K = 0$, and $\Delta S = 0$—nothing changed—and resources were consumed. Noise is not rest, stillness, or recovery. Rest produces ΔS (restoration, neurological repair, emotional regulation). Noise is activity that mimics the appearance of signal while producing no state change.

Noise decomposes into four measurable components:

Redundancy (R). Repeated processing of the same information or task without incremental value. Re-reading the same email a third time. An agent retrying a failed API call with identical parameters. A model regenerating output that is structurally identical to its previous generation. A meeting that covers ground already covered in the previous meeting. Redundancy is measurable by comparing input-output pairs across iterations: if the delta between iteration n and iteration n+1 approaches zero, the time spent on n+1 is redundancy. Measured in hours of repeated processing.

Friction (F). Time consumed by unnecessary intermediation, latency, bureaucracy, or poor system design that adds no value to the outcome. Waiting in authentication queues. Navigating a seven-step approval process for a decision that requires one step. Format conversions that exist because two systems were never integrated. Cold starts on infrastructure that should be warm. Friction is measurable as the difference between the time a process actually takes and the time the value-producing steps within it take. Measured in hours of overhead.

Drift (D). Time spent moving away from an intended objective without course correction. Scope creep in a project that extends timelines without extending value. An agent pursuing a dead-end reasoning chain for forty seconds before backtracking. A person falling into a content rabbit hole that started as research and ended as aimless browsing. A company spending eighteen months on a strategic pivot that never completes. Drift is measurable by tracking the cosine similarity between activity trajectory and stated objective over time—when the trajectory diverges and does not reconverge, the divergent time is drift. Measured in hours of misdirected effort.

Dissipation (Dp). Energy—cognitive, computational, financial—consumed without producing any of the three signal dimensions. GPU cycles on inference that gets discarded by the user. Human attention spent on content that leaves no memory, no skill, no changed state. Processing power consumed by orchestration overhead that does not improve the final output. Marketing spend that generates impressions but no conversions, enquiries, or brand state change. Dissipation is measurable as total resources consumed minus resources that demonstrably contributed to ΔO, ΔK, or ΔS. Measured in hours of wasted capacity.

The total noise produced by any system or process is:

$$\text{Noise (N)} = R + F + D + Dp$$

Each component is denominated in hours, making noise directly comparable to signal and the ratio between them dimensionless.

3.4 The Temporal Signal-to-Noise Ratio (TSNR)

For any system, agent, platform, company, or process:

$$\text{TSNR} = (\Delta O + \Delta K + \Delta S) / (R + F + D + Dp)$$

The structure of this ratio follows a lineage that spans eight decades of measurement science. Claude Shannon formalised the Signal-to-Noise Ratio at Bell Labs in the 1940s to measure the quality of communication channels: how much of a transmitted signal is information versus how much is interference. The concept was adopted into electrical engineering, then audio processing, then medical imaging—where SNR became the defining quality metric for MRI scanners, determining whether a scan reveals a tumour or shows only static. More recently, quantitative finance adopted SNR to separate genuine alpha from market noise.

In each domain, the mathematical structure is identical: a ratio of useful output to wasted output. The units change—voltage, photon counts, basis points, hours—but the architecture does not. TSNR extends this architecture into temporal economics with one important distinction. In classical SNR applications, signal and noise are measured through physical instrumentation: voltage meters, photon detectors, frequency-domain analysers. In TSNR, signal and noise are classified through the state-change test and the four noise components (R, F, D, Dp), which involve structured estimation rather than instrument readings. The mathematical form is preserved. The measurement domain is new. Where an MRI scanner asks "how much of this image is diagnostic information versus thermal interference," TSNR asks "how much of this system's time expenditure produces state change versus how much spins without result." The question is the same. The subject has changed.

A TSNR greater than 1 indicates a net signal-positive system—it produces more state change than waste. A TSNR less than 1 indicates a system that consumes more time in noise than it generates in signal. A TSNR of exactly 1 means the system breaks even: for every hour of state change it produces, it wastes an hour.

This ratio is universally applicable:

Domain	Signal Components	Noise Components	Typical TSNR Range
AI models	ΔO: usable outputs generated ΔK: information synthesised ΔS: decisions enabled	R: regenerated/discarded outputs F: latency, retries D: off-target responses Dp: compute on unused output	2.0 – 15.0
Autonomous agents	ΔO: tasks completed ΔK: data gathered for decisions ΔS: system states transitioned	R: repeated tool calls F: orchestration overhead D: wrong strategy pursuit Dp: failed reasoning chains	3.0 – 20.0
Businesses	ΔO: products shipped, deals closed ΔK:	R: meetings rehashing decisions F: approval	0.5 – 5.0

	institutional learning ΔS: operational improvements	bureaucracy D: incomplete pivots Dp: spend without conversion	
Individuals	ΔO: work completed ΔK: skills acquired ΔS: health, relationships improved	R: re-consuming same content F: commute, admin overhead D: started-not-finished tasks Dp: attention without retention	0.3 – 3.0
Infrastructure	ΔO: throughput delivered ΔK: operational data captured ΔS: self-healing, auto-scaling	R: health checks without action F: cold starts, warm-up cycles D: load balancer oscillation Dp: idle resource consumption	5.0 – 50.0

The Jobs principle maps precisely onto this framework. His signal-to-noise ratio was high because he operated with minimal redundancy (decisions were made once), minimal friction (unnecessary process was eliminated), minimal drift (the objective was never out of focus), and minimal dissipation (every unit of energy in the room was directed at the problem). That is a TSNR description, not a biographical anecdote.

3.5 Noise in AI Systems: The Slop Problem

The current AI landscape presents a specific problem. AI systems are simultaneously the most powerful time-saving tools ever built and the most prolific generators of temporal noise in history. Generative AI produces vast quantities of low-quality content that consumes human attention without producing state change. Articles that contain no information the reader did not already possess. Images that communicate nothing the viewer retains. Code that must be debugged and rewritten at a time cost approaching the cost of writing it from scratch. Emails that require as much effort to parse as they saved to compose.

Each of these represents a negative time transfer: the system consumed computational time to produce output that then consumes human time to evaluate, filter, or discard. The TSNR captures this directly. A code generation tool that produces working code 70% of the time and broken code 30% of the time has a measurable TSNR that accounts for the debugging hours imposed by the broken 30%. A content platform whose output requires extensive human editing to reach usable quality may have a TSNR below 1—it consumes more human time than it saves.

This is not an argument against AI. It is an argument for measuring AI with the correct instrument. The tools that produce genuine state change will score well. The tools that generate volume while shifting the time cost to the consumer will be exposed by a metric that cannot be gamed by output quantity alone.

3.6 The Consumption-Productivity Index (CPI-T)

Building on the TSNR, we introduce a composite metric for evaluating the net temporal character of any system:

$$\text{CPI-T} = (T\_signal - T\_noise) / T\_total$$

Where:

T_signal = Total time producing state change (ΔO + ΔK + ΔS)
T_noise = Total time consumed by noise (R + F + D + Dp)
T_total = Total time engaged with the system

CPI-T ranges from -1 (pure temporal consumption—all engaged time is noise) to +1 (pure temporal productivity—all engaged time produces state change). A CPI-T of 0 indicates temporal neutrality.

System	CPI-T	Primary Signal	Primary Noise
Autonomous scheduling agent	+0.87	ΔO: meetings booked, conflicts resolved	F: calendar API latency
AI diagnostic tool (medicine)	+0.91	ΔS: patient states assessed	R: redundant differential checks
AI code assistant (senior dev)	+0.68	ΔO: functioning code produced	Dp: discarded suggestions
AI code assistant (junior dev)	+0.25	ΔK: patterns learned from output	D: debugging misdirected code
AI content generator (marketing)	+0.18	ΔO: publishable drafts	R: near-identical variations
Algorithmic social feed	-0.45	Minimal ΔK on rare occasions	D: rabbit holes; Dp: attention drain
Excessive push notifications	-0.60	Near-zero across all dimensions	F: context-switch cost; Dp: attention

CPI-T and TSNR together provide a complete picture of temporal character. TSNR tells you the ratio of productive to wasted time. CPI-T tells you the net direction—whether the system, on balance, creates time or destroys it.

Part IV: The TIME Score

4.1 A Universal Rating

The TIME Score is the synthesis of every mathematical layer described above. It is a single number, scored from 0 to 100, that represents the net time value of any AI system, agent, platform, company, or process. It is designed to be immediately legible to a consumer, an investor, or a policymaker, while being grounded in auditable mathematics.

The TIME Score is composed of four weighted components:

$$\text{TIME Score} = (w1 \times QATU\_n) + (w2 \times TSNR\_n) + (w3 \times CPIT\_n) + (w4 \times TLR\_n)$$

Where:

QATU_n = Normalised Quality-Adjusted Time Units (raw time liberation, 0–100)
TSNR_n = Normalised Temporal Signal-to-Noise Ratio (signal quality, 0–100)
CPIT_n = Normalised Consumption-Productivity Index (net temporal direction, 0–100)
TLR_n = Normalised Time Leverage Ratio (computational efficiency, 0–100)

Default weights: w1 = 0.35, w2 = 0.25, w3 = 0.25, w4 = 0.15. These reflect the primacy of actual time liberated (QATU), the importance of signal purity (TSNR), the net directional impact (CPI-T), and the computational efficiency of the liberation (TLR). Weights may be adjusted for sector-specific applications.

The TIME Score is the number that appears on the leaderboard, on the badge on a company’s website, in an investor’s portfolio analysis, and on a consumer’s personal dashboard. Everything in Parts II and III exists to make this single number defensible.

4.2 Negative TIME Scores

Systems with a net-negative temporal impact—those that consume more human time than they liberate through addictive design, noise generation, or attention exploitation—receive a TIME Score below 50, with severely extractive systems scoring in single digits.

The existence of negative scores is what separates a standard from a marketing exercise. A framework that only measures upside is promotion. A framework that measures both sides of the ledger is accounting.

4.3 The TIME Score for Individuals

The TIME Score applies to people as well as systems. By aggregating the TIME Scores of all AI tools, platforms, and agents a person uses—weighted by their usage patterns—we produce a Personal TIME Score that answers a direct question: is technology, on balance, giving you time or taking it?

A person who uses high-TSNR autonomous agents and avoids low-CPI-T attention platforms will carry a high Personal TIME Score. A person trapped in engagement loops with occasional AI assistance will carry a low one. The score is a mirror. It does not judge. It measures.

Part V: The Economics of Time

5.1 Time as the Base Layer

We arrive at the central claim of this paper: time is not a derivative of value. Time is value itself.

Every other economic quantity can be derived from time. The price of a good is the claim it makes on someone else's time. The value of a skill is determined by how much time it takes to acquire and how much time it saves others. The worth of a company is a function of the future time its products will command or liberate. GDP is an aggregate measure of how a nation's population spent its time during a given period.

Money is a representation of time. It always has been. But the representation was so convenient, so ubiquitous, and so embedded in institutional infrastructure that we forgot what it represented. We began optimising for the representation instead of the underlying reality. We built an economy that maximises monetary throughput rather than temporal liberation.

The autonomous intelligence revolution strips this illusion away. When machines produce economic output without human time input, the money continues to flow but the time equation changes. The companies that will define the next era are not those that generate the most revenue. They are those that liberate the most time—and the TIME standard is the instrument that makes this visible.

5.2 Time Markets

If time can be measured, it can be priced. If it can be priced, it can be traded. This is not speculation—it is the consequence of standardisation.

Consider what happened when carbon emissions were standardised and measured. Within a decade, carbon credits became a tradeable market worth hundreds of billions. The act of measurement created the market. The same dynamic applies to time.

Once TIME Scores and QATU outputs are standardised and trusted, market mechanisms follow. Companies could issue instruments whose returns are tied to QATU targets rather than revenue. Index funds could weight holdings by TIME Score rather than market capitalisation. Early-stage investments could vest based on time impact milestones. The specific instruments are secondary. The point is that measurement precedes markets, and markets follow measurement as reliably as trade followed the standardisation of weights and measures.

5.3 Time and Competition

When time becomes the measure, competition transforms. Companies no longer compete solely on features, price, or growth. They compete on time impact. The question shifts from “which product is cheapest” or “which has the most features” to “which gives me the most time back?”

This reframes the entire technology industry. A search engine that delivers an answer in 3 seconds competes on time against one that delivers it in 30. An agent that completes a task autonomously competes on time against one that requires supervision. A platform that returns an hour of your day competes on time against one that consumes an hour through engagement design.

Time-based competition forces honesty. You cannot fake time savings. You can inflate revenue with aggressive sales tactics. You can inflate user counts with growth hacking. But you cannot inflate the actual hours returned to actual human beings. The TIME Score resists manipulation because its input—whether a human got their time back—is a physical fact, not an accounting construct.

5.4 Time at Civilisational Scale

If every AI system, every agent, every platform is scored on time impact, we gain a metric that no civilisation has possessed: a real-time measure of whether technology is serving humanity or extracting from it.

Aggregate TIME Scores across an entire economy answer something GDP cannot: are the citizens of this society getting their time back? Are they spending more hours on things that produce state change—output, knowledge, improved conditions? Or are they losing hours to redundancy, friction, drift, and dissipation?

A nation where aggregate TIME Scores are rising is a nation where technology is fulfilling its promise. A nation where they are falling is a nation where technology has been captured by interests that extract time rather than return it. This is a more honest measure of progress than any monetary aggregate, because it measures the thing that matters to human beings: whether they have enough time to live well.

Time spent on algorithmic feeds at population scale is not a neutral fact. It is a measurable, scoreable consumption of human potential. The TIME standard does not moralise about this. It counts it. And counting, done consistently, changes everything.

5.5 The Time-Energy Entanglement

Time and energy are not independent variables. They are coupled at every level of economic activity. Every hour of human labour consumes caloric energy. Every hour of computation consumes electrical energy. Every hour of manufacturing consumes mechanical and thermal energy. Every hour of transport consumes kinetic energy. There is no economic time expenditure that does not have a corresponding energy expenditure, and there is no energy expenditure that does not occur across a duration of time.

This entanglement means that the TIME standard, by measuring and optimising the time dimension, simultaneously acts on the energy dimension. The four components of temporal noise—redundancy, friction, drift, and dissipation—each carry direct energy costs:

Redundancy (R) consumes energy by performing the same computation, the same meeting, the same process multiple times without incremental result. Every redundant GPU cycle is wasted electricity. Every redundant business meeting is wasted human caloric and cognitive energy. Every redundant manufacturing run is wasted material and thermal energy.

Friction (F) consumes energy through intermediation that adds no value. Cold starts on servers that should be warm waste electricity. Commutes that could be eliminated waste fuel. Approval chains that add no information waste the metabolic energy of every person in the chain.

Drift (D) consumes energy by moving in the wrong direction before correcting. An agent that pursues a dead-end reasoning chain for thirty seconds before backtracking has consumed GPU energy on a path that produced nothing. A company that spends eighteen months on an incomplete strategic pivot has consumed the energy of every employee involved.

Dissipation (Dp) is energy loss by definition. It is the direct translation of the thermodynamic concept into economic terms: energy consumed that does not contribute to any state change in the system.

A system with a high TSNR is therefore, by construction, more energy-efficient than a system with a low TSNR. Reducing temporal noise reduces energy waste. They are the same optimisation viewed from two different measurement frames. QATU-based outcome analysis applied to any process simultaneously reveals its time efficiency and its energy efficiency, because you cannot improve one without improving the other.

This entanglement has a practical implication for resource management that extends far beyond AI. Every industry that adopts TIME scoring gains, as a side effect, a granular map of where energy is being wasted. A logistics company that reduces its temporal friction (unnecessary routing delays, redundant dispatching, drifted delivery sequences) simultaneously reduces its fuel consumption. A hospital that reduces its temporal redundancy (repeated diagnostic tests, duplicated intake procedures, friction in referral chains) simultaneously reduces its electricity, staffing energy, and equipment wear. The TIME standard is, without any modification, also an energy audit.

Case Study: Content Production

The economics of large-scale content production illustrate the time-energy entanglement at industrial scale. A major streaming platform reportedly spends in the range of $17 billion annually on content, producing several hundred original titles per year. Each production involves months of development, weeks of physical shooting with crews of hundreds, months of postproduction, and substantial physical infrastructure: studio facilities, lighting, transport, catering, location logistics, and global distribution of physical equipment and personnel.

The energy footprint of a single prestige production is substantial. Studio electricity for lighting and climate control across a multi-month shoot. Fuel for transporting cast, crew, and equipment across locations. Server farm energy for rendering, editing, colour grading, and visual effects processing. The aggregate energy cost of a production slate numbering in the hundreds of titles is measured in hundreds of millions of kilowatt-hours annually.

The hit rate on this investment is structurally low. Industry data suggests that a minority of original productions achieve viewership targets that justify their cost. The remainder represent drift and dissipation at scale: time and energy spent producing content that does not achieve its intended state change (audience engagement, subscriber retention, cultural impact).

Now consider the same output produced through generative AI. An AI-generated feature film consumes GPU compute time measured in hundreds or low thousands of hours—a fraction of the human production timeline. It requires no physical sets, no location transport, no catering, no studio electricity for lighting rigs. Its energy footprint is concentrated entirely in computation, which is both measurable and declining in cost per unit as hardware improves.

The QATU analysis is direct. If 100 productions consume 500,000 aggregate human hours and $1.7 billion in cost, and 30 of those productions achieve their intended audience impact, then 70% of the time and energy expenditure was noise—drift (pursuing creative directions that did not connect) and dissipation (resources consumed on content that produced no measurable state change in the audience). The TSNR of the traditional production model might be approximately 0.43. An AI production pipeline that can test, iterate, and produce at a fraction of the time and energy cost—and can produce a far greater volume of variants to discover what resonates—would have a categorically higher TSNR, even if individual AI productions are lower in craft quality.

The volume advantage compounds the efficiency gain. Where a traditional pipeline produces hundreds of titles and hopes a sufficient fraction succeed, an AI pipeline could produce thousands of variants, test audience response in hours rather than months, and allocate full production resources only to the concepts that demonstrate traction. The noise in the system—the drift and dissipation of producing expensive content that does not find an audience—is reduced by orders of magnitude. And because time and energy are entangled, the energy savings follow the time savings directly.

This pattern generalises across every industry where production involves physical infrastructure, human coordination at scale, and uncertain hit rates. Pharmaceuticals. Architecture. Product design. Advertising. Education content. In each case, QATU-based analysis reveals the time-energy waste hidden in traditional production models, and the time-energy efficiency available through autonomous alternatives. The TIME standard does not prescribe which approach to use. It measures the difference, and the measurement speaks for itself.

Part VI: The World That Comes Next

6.1 Buying and Selling Time

We are approaching a world in which time is directly transactable. Not as a metaphor but as a structure. When a person subscribes to an autonomous agent that handles their administrative work, they are purchasing time. When an investor funds a company with a high TIME Score, they are investing in time production. When a government deploys AI systems across public services and measures the QATU output, it is allocating resources toward time liberation.

The TIME standard makes these transactions visible and comparable. It allows a consumer to evaluate any AI tool not by its feature list but by its time return. It allows an investor to evaluate any AI company not by its revenue trajectory but by its time production. It allows a society to evaluate its technological infrastructure not by its economic output but by whether it is giving citizens their lives back.

6.2 The Competition for Human Flourishing

When time is the measure, something happens to incentives. Companies that want high TIME Scores must actually save people time. They cannot game the metric through engagement tricks because those tricks increase noise and reduce the score. They cannot inflate it through marketing because the score is based on state changes, not claims. The only path to a high TIME Score is to build something that gives people hours back, does so with high signal, low noise, and fair distribution.

This creates a competitive dynamic oriented around human flourishing rather than extraction. The most valuable companies in a time-measured economy are those that liberate the most time at the highest quality for the most people. Time puts humanity in check. It forces the question that money alone never asks: did this technology make life better, or did it just make money?

6.3 The Inelastic Asset

Time has a property that no other economic asset possesses: perfectly inelastic supply at the individual level. Every human being receives exactly 24 hours per day, regardless of wealth, status, or demand. No amount of capital can purchase a 25th hour. No technology can create one.

What autonomous intelligence can do—and what the TIME standard measures—is liberate hours that were previously locked in labour, administration, waiting, and noise. This is not the creation of new time. It is the unlocking of existing time. The economics of time liberation are closer to the economics of unlocking trapped capital than to the economics of production.

The total addressable market for time liberation is all human time currently spent on tasks that autonomous systems could perform. Conservative estimates place this at 60–70% of professional working hours in developed economies. At a global scale, the time liberation opportunity represents tens of billions of hours annually—an economic quantity larger than any market currently measured, invisible because we lack the instrument to see it.

6.4 TIME-Informed Currencies

If time is the base layer of value, and QATUs are its unit of measure, a question follows: can QATUs inform or underpin new forms of currency?

The Bretton Woods system pegged currencies to gold. Each dollar represented a claim on a physical quantity of a scarce element. That system collapsed in 1971, but the principle it embodied was sound: currencies function best when anchored to something real. The post-Bretton Woods era of fiat currency severed that anchor, allowing monetary supply to float against political and market forces rather than physical constraints. Cryptocurrencies attempted to restore scarcity through computational proof-of-work, but anchored value to energy expenditure rather than to any measure of human utility.

QATUs offer a different anchor: verified time liberation. A QATU-informed currency would derive its value not from scarcity of a metal, not from difficulty of a computation, but from the measured, audited fact that human time was returned. Each unit would represent a claim backed by a real-world state change—hours that actual humans no longer need to spend on tasks that machines now perform.

This is not a proposal to issue a token. It is an observation about what becomes possible once the measurement layer exists. The sequence is: standardise the unit, build the audit infrastructure, accumulate trusted data, and then allow the market to determine whether QATU-denominated instruments emerge. The standard does not need to prescribe the financial architecture. It needs to provide the measurement that any financial architecture would require.

Within closed ecosystems—agent marketplaces, autonomous service platforms, inter-agent settlement layers—QATUs have an immediate use as a unit of account. When Agent A requests a sub-task from Agent B, the natural unit of exchange between them is the time impact of the transaction on the human principal. Dollars are a translation layer that agents do not natively understand. Time is the dimension in which they operate. QATU-denominated settlement between agents is not a speculative future. It is a more natural expression of the value being exchanged than any fiat denomination.

Whether this evolves into a broader currency, a settlement protocol, a staking mechanism, or something without a current analogue is a question for markets and institutions to answer. The TIME standard provides the foundation. What is built on that foundation will be determined by the same force that has shaped every financial innovation in history: the need to measure and exchange something that people recognise as valuable.

6.5 Acceleration

The promise of the TIME standard is the acceleration of civilisation itself. If we can measure how much time technology returns to humanity, we can optimise for it. If we can score the time impact of every system, we can direct resources toward the ones that matter. If we can make time visible as an economic quantity, we can build an economy that treats human time as what it is: the only non-renewable resource that every person on earth shares equally.

The tools exist. The autonomous agents are running. The intelligence is deployed. The only thing missing is the measurement.

This paper provides it.

Time is value. Measure it.

Appendix: Technical Summary

Term	Definition	Unit
T	Raw time liberated: H - A	Hours
TDM	Time Delta Metric: $\Sigma(H_i \times U_i \times D_i)$	Weighted hours
QATU	Quality-Adjusted Time Unit	1 hour at U=1, D=1
CTS	Computational Time Stack	Normalised compute-hours
TLR	Time Leverage Ratio: T / CTS	Dimensionless ratio
TCR	Time Compounding Rate	Dimensionless ratio
E	Temporal Entropy	Bits (normalised 0–1)
eQATU	Entropy-adjusted QATU	Adjusted hours
ΔO	Delta Output (signal)	Hours producing artefacts
ΔK	Delta Knowledge (signal)	Hours producing learning
ΔS	Delta State (signal)	Hours producing transitions
R	Redundancy (noise)	Hours of repeated processing
F	Friction (noise)	Hours of overhead

D	Drift (noise)	Hours of misdirected effort
Dp	Dissipation (noise)	Hours of wasted capacity
TSNR	Temporal Signal-to-Noise Ratio	Dimensionless ratio
CPI-T	Consumption-Productivity Index	Range: -1 to +1
TIME Score	Composite rating	0–100
H_relative (t)	Rolling baseline: A_best(t-1)	Hours
E_cost	Energy cost entangled with temporal noise	kWh (correlated)

—

TIME

Ash Koosha, 2025–2026