Supervised Automation

The New Paradigm of Outcome-based Productivity

Whitepaper

Applied AI

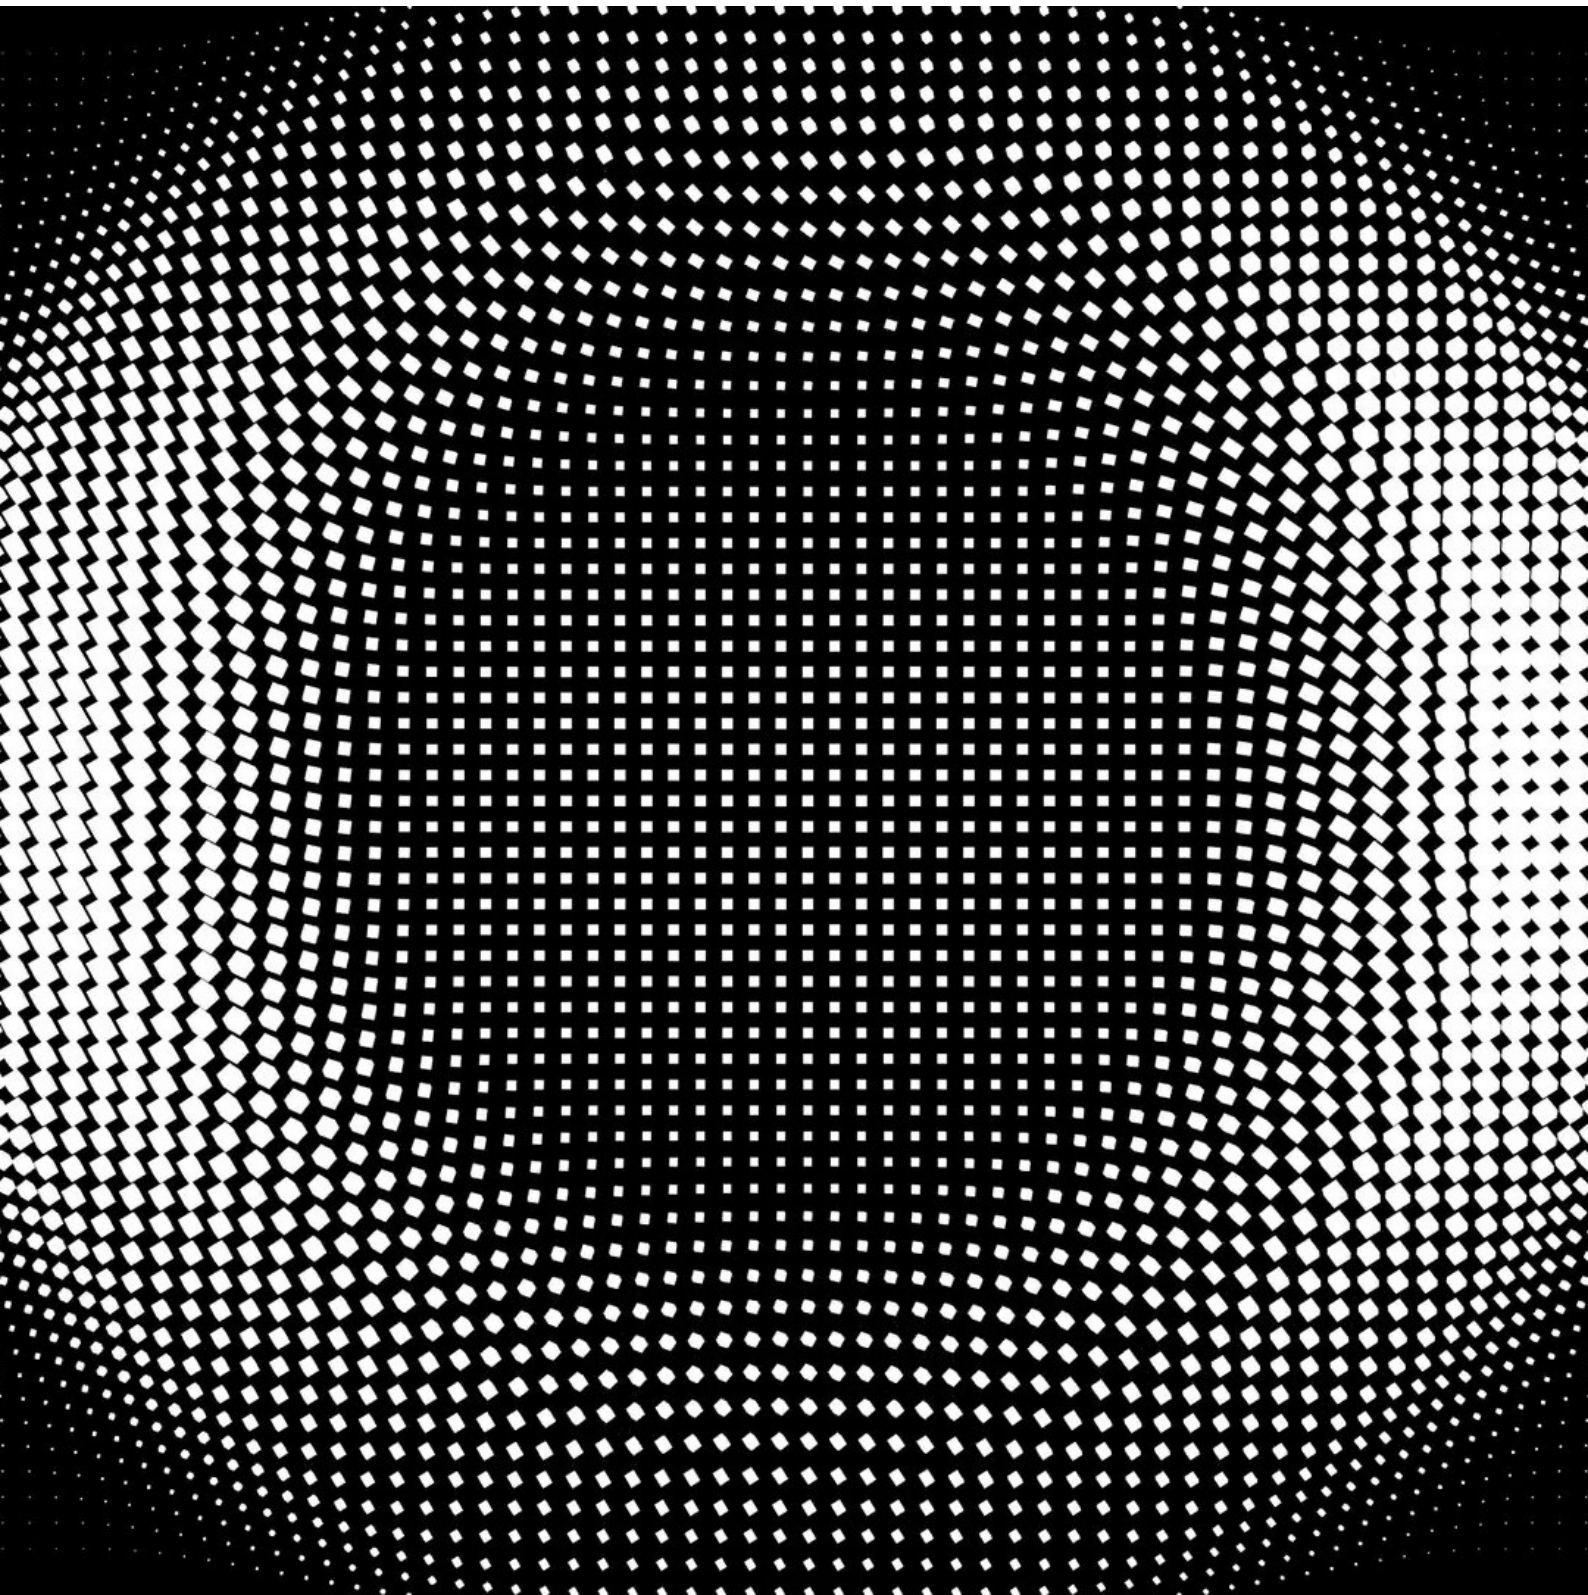

A Letter from the CEO

We stand at the dawn of a new era—a time when the boundaries between man and machine blur, and the possibilities for human advancement are boundless. At AppliedAI, our mission is clear:

To prepare humanity for a post-AI world where work remains a source of meaning, fulfillment, and endless potential.

Work is more than a means to an end; it is the heartbeat of our existence, the canvas upon which we paint our aspirations and the legacy we leave behind. Yet today, the essence of work is under unprecedented threat. The employment frameworks we've long relied upon are becoming obsolete—relics of a bygone era.

This concern isn't just economic; it's deeply human. Knowledge work, representing about one-third of global work expenditure, is particularly susceptible to dislocation. With 1 billion knowledge workers earning $4 trillion annually, the stakes are high.

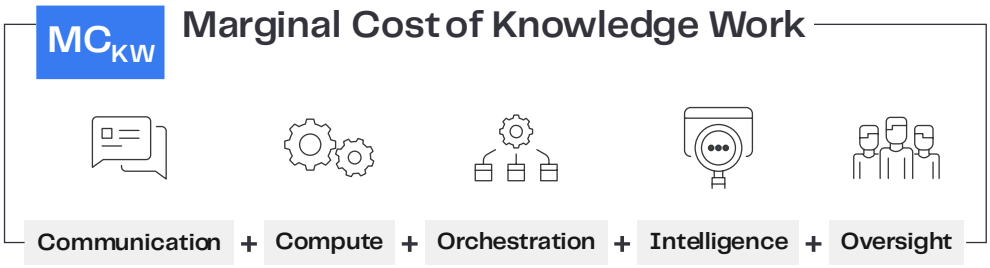

Until now, the cost of intelligence and the cost of audit were bottlenecks to progress. Knowledge work logistics has been improving through robotic process automation (RPA), enterprise resource planning (ERP), and SaaS platforms. Intelligence costs, due to recent advances in foundational models, compute power, and application-layer systems are increasingly available on-demand and at scale. However, the cost of audit—the tech-resistant need for human oversight and quality control—remains and shall remain. In fact, it may grow as governments and shareholders increase oversight on AI automation even as AI adoption continues to rise.

We believe that labor is an integral part of the future solution. AI is not a substitute good but a complement. That's why we've toiled to build **Opus**, our solution to reimagine the future of work. We provide not just task-level automation, but the whole job-to-be-done—the finished work that clients can trust and put into use. In the context of rapid and seismic change in the field of AI, we have chosen to focus on what will not change: people will want their work done more cheaply, faster, and with greater quality.

Opus will catalyze a new labor structure for knowledge work that will lead to a 10× reduction in marginal cost. Imagine a world where the burdens of mundane administration are lifted—a world where brilliant ideas no longer die on the vine due to bureaucratic inertia. Many a Beethoven lived and died before the advent of the piano, their unique poetry unheard because there was no mechanism to showcase it.

Opus aims to eliminate these soul-crushing tasks, allowing labor to focus on the edge cases and the creative.

Will this lead to the annihilation of knowledge workers? Absolutely not, nor is it designed to. We believe that AI will increase human productivity tenfold, while simultaneously expanding the opportunities even more. Once we remove the barriers clogging the fountain of human innovation, even if humans perform only quality control, there will be more than enough meaningful work for everyone.

We believe that a vast segment of the economy that the SaaS revolution has overlooked—services—can now be automated, giving rise to services-as-software. Many legacy industries have evolved into on-demand models—transportation, finance, accommodation, healthcare, and education. It's time for knowledge work to do the same.

Opus stands unique in the AI industry in two dimensions:

Humanist and Optimistic

Seeking human augmentation, not human replacement.

Liberating and empowering

Unleashing humans with greater flexibility, income, and control.

We're building a home for the Ronin knowledge worker and offering a new path for sunsetting BPOs. The last major industrial revolution challenged labor, leading to the alienation of workers from the product of their work and surplus profits not accruing to them—resulting in significant geopolitical implications. In this age of intelligence, we have the opportunity—and the responsibility—to do better. By arming labor with the tools to compete and thrive in a post-AI world, we can foster a future where prosperity is shared, and innovation knows no bounds.

We believe in a future where technology amplifies human potential, where work is redefined for the betterment of all, and where innovation is unshackled from the chains of the mundane.

With conviction & gratitude,
Arya H. Bolurfrushan

Executive Summary

Imagine a world where work isn't measured in hours but in impact. Where AI doesn't replace human ingenuity – it amplifies it. Where the relentless churn of emails, meetings, and administrative tasks is not a necessary evil but a relic of the past. This isn't a utopian fantasy or the fevered dream of a Silicon Valley evangelist. **It's happening now.**

For the past decade, AI has flirted with productivity gains – automating customer service chats, crunching data, and even attempting (and sometimes failing) to mimic human creativity. But the true AI revolution is only just beginning. The real transformation isn't simply automating the mundane; it's **redefining how work is structured, valued, and executed.**

Consider this: a huge portion of "knowledge work" – up to 50% by some estimates – is already structured around repeatable workflows. Yet organizations still compensate workers based on time spent rather than outcomes achieved. Businesses pay for hours, not insights; for presence, not performance. This outdated model has led to stagnation, where productivity plateaus and knowledge workers spend more time managing work than doing it. In essence, we've been measuring and rewarding the input (hours) instead of the output (results). But what if AI could break this cycle? What if we could pay for results instead of hours and let people reclaim the time currently lost to drudgery?

Enter **Opus**, the world's first Knowledge Work AI platform.

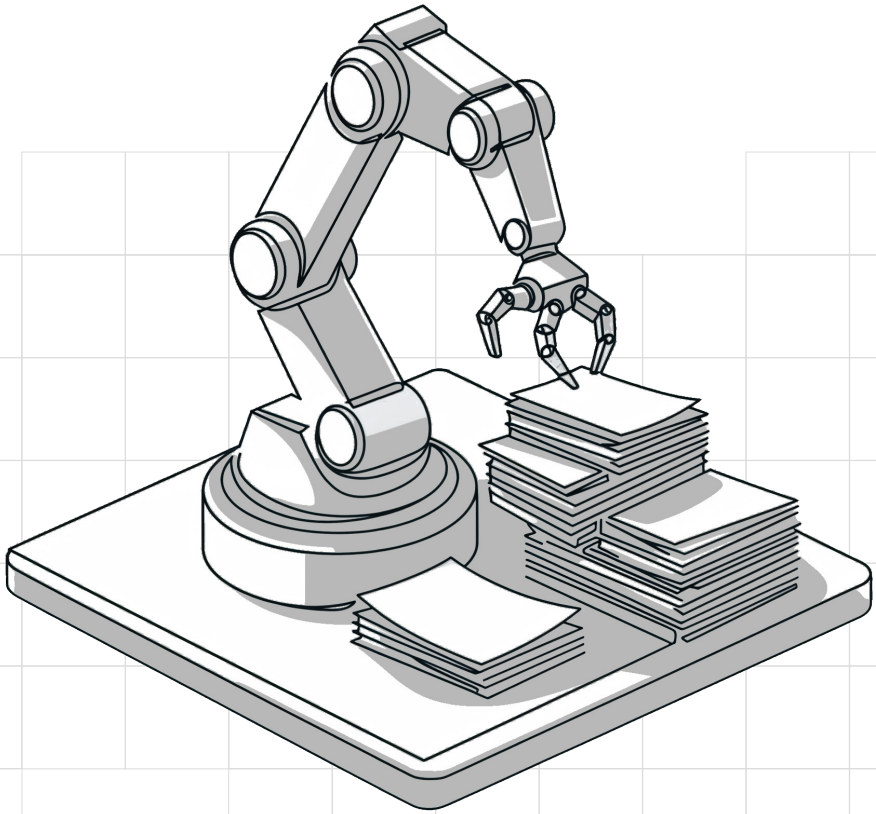

It's not just another automation tool – it pioneers a new paradigm called **Supervised Automation**, a human-in-the-loop model that fundamentally shifts how businesses operate. Opus leverages AI to create fluid, outcome-based workflows that liberate human talent from repetitive tasks. Instead of traditional outsourcing, organizations can tap into a dynamic, cloud-based workforce where AI handles the heavy lifting and humans serve as expert reviewers or "judges" of the output. In this model, the workforce isn't a static roster of employees or an offshore team – it's a blend of AI and human expertise distributed globally, compensated not for showing up, but for delivering results.

The impact? A system that is more equitable, more efficient, and ultimately more human. When routine work is automated and only the exceptions or creative tasks require human input, people can focus on what they do best. The value of work shifts to quality and outcomes. A finance analyst no longer spends 4 days compiling reports (because an AI does it in minutes) but spends that time devising strategy; A customer support agent now supervises an army of AI bots handling inquiries, stepping in only for the most complex, human situations. This is the promise of Supervised Automation – **boosting productivity while enhancing human roles, not diminishing them.**

The question is no longer whether AI will change the way we work – that debate is settled. The real question is how we will embrace this future: will we design technology to enhance human potential or to replace it? The pages ahead lay out how AI, when designed for collaboration rather than competition, can unlock a new era of productivity, purpose, and prosperity. We will explore the economic trends that led here (from the outsourcing boom to the rise of AI), the scientific breakthroughs making it possible (from knowledge graphs to large work models), and the human-centered approach required to get it right (keeping experts in the loop, ensuring transparency and ethics).

The future of work has arrived.

Turn the page.

See how it all comes together.

The Evolving Nature of Work

1.1 Productivity, Technology, & Human Purpose

Work has continually evolved – from the farms of the Agricultural Revolution to the factories of the Industrial Revolution, and now to the digital offices of the Information Age. Each transformation boosted economic output and living standards by making labor more productive. Productivity, simply put, measures how efficiently we turn inputs (labor and capital) into outputs (goods, services, ideas). It's the engine of prosperity. Over decades, a society's ability to improve its standard of living depends on its ability to raise output per worker. As Nobel laureate Paul Krugman famously noted,

"Productivity isn't everything, but in the long run it is almost *everything*".

Peter Drucker coined the term "knowledge work" in 1959 to describe labor that uses knowledge, creativity, and decision-making to create value. While technological advancements dramatically increased productivity in manufacturing and services in the 20th century, measuring productivity in modern knowledge work is more complex. Unlike making widgets on an assembly line, the outputs of knowledge work (insights, designs, strategies) are intangible and hard to quantify. In recent decades, we've seen "computers everywhere except in the productivity statistics," a phenomenon often referred to as the Solow Paradox (after Robert Solow's 1980s observation).

While digital tools became ubiquitous, measured productivity growth in many sectors stagnated. Knowledge workers were equipped with email, spreadsheets, and software, yet often they were doing more work about work – coordinating, communicating, and administrating – rather than significantly boosting output. The explosion of information and communication added as much overhead as it removed. In the United States, for example, labor productivity growth averaged only ~1.3% per year from 2005–2016, less than half the 2.8% annual growth of 1995–2004, despite the proliferation of software and internet tools.

GDP per capita (Geary–Khamis dollars)

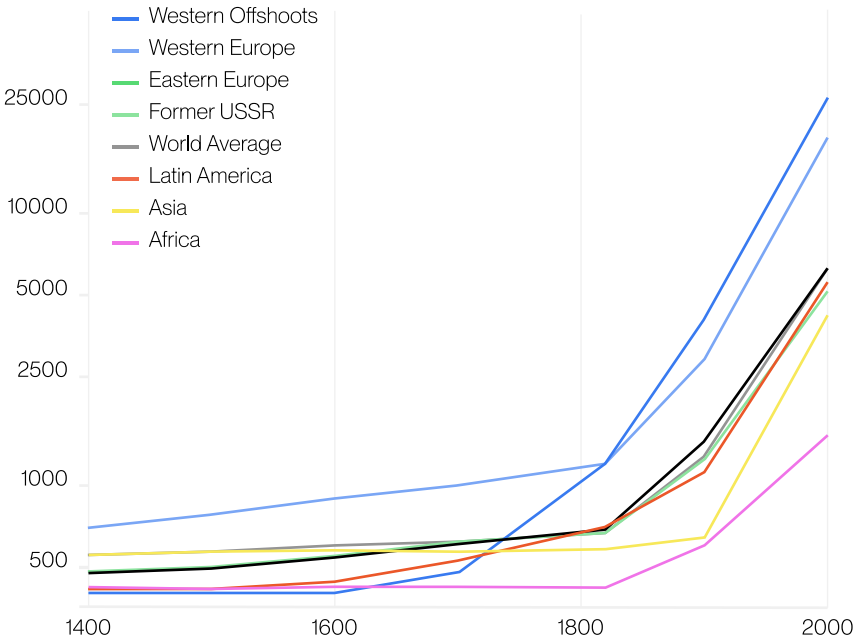

"GDP per capita (Geary-Khamis dollars). Global GDP per capita over time, showing centuries of stagnation followed by an exponential rise after industrialization. Economic growth and human progress have always been intertwined with technology. For centuries, world GDP per capita remained nearly flat; then the industrial and digital revolutions led to exponential growth in living standards. Total Factor Productivity (TFP), which captures the portion of output not explained by labor or capital alone (essentially the contribution of technology and efficiency), surged with the general-purpose innovations of the 20th century. Higher productivity means more output per worker, enabling higher incomes and living standards.*

Historic world GDP per capita. svg - Wikimedia Commons

This "productivity paradox" suggests that the last wave of digital transformation did not translate into the productivity gains many expected. Software often simply digitizes existing processes without structurally reinventing them.

Yet, work is not just about output and metrics; it is a profoundly human endeavor that gives us purpose and identity. People seek meaningful work – roles that leverage their creativity, skills, and passion. In its best form, work is a source of achievement and fulfillment, not just a paycheck. The tragedy of today's work life is that highly skilled individuals often spend a large portion of their day on routine tasks that don't tap into their true talents. The challenge – and opportunity – for technology is to free humans from this drudgery so that work can be refocused on what humans find meaningful.

Importantly, history shows that as technology advances, the nature of work continually adapts. Consider the evolution of medical records management over the past decades. As recently as the early 2000s, thousands of medical records clerks maintained vast libraries of paper charts, meticulously filing each document and retrieving records for appointments. Today, that paper-based system has all but vanished due to Electronic Health Record (EHR) systems. Yet those workers were not left permanently unemployed; their roles evolved. Many became EHR specialists, data integrity analysts, and compliance coordinators. Their productivity soared—a single EHR specialist now supports thousands of patient records that would have required multiple clerks to manage physically, while simultaneously improving record accuracy and enabling critical insights through analytics.

Similarly, financial analysis once involved rooms of professionals hunched over spreadsheets, manually gathering and formatting data for reports. By the 2010s, automated data pipelines and visualization tools had transformed this landscape, as analysts began to focus on insights rather than data preparation. Again, jobs evolved: the analyst role shifted toward strategic advisory functions and scenario modeling, with individual analysts now delivering recommendations that drive millions in business value instead of simply reporting historical figures. Former spreadsheet masters transitioned into data science, business intelligence, and strategic finance roles, where they could analyze 10-20 times more data while providing more timely, accurate, and actionable insights to leadership.

Automation frees humans from routine drudgery to take on new, often more complex tasks. When mechanized looms appeared, weavers feared for their livelihoods – yet the textile industry ultimately employed more people in higher-skilled roles managing machines. When ATMs spread in the 2000s, bank tellers were expected to disappear – instead, teller jobs shifted toward relationship banking, and the number of bank branches grew, employing even more staff in advisory capacities. The pattern is clear:

As technology handles "busy work," humans pivot to roles that leverage uniquely human strengths – creativity, problem-solving, and interpersonal skills.

These historical transformations underscore a key point: **automation changes jobs, but it doesn't mean the end of work.** Supervised Automation accelerates this shift, aiming to relieve workers of structured, repetitive tasks and enable them to concentrate on strategic, creative, and judgment-intensive activities. Organizations can nourish innovation and productivity, rather than just feeding the never-ending appetite of routine processes.

1.2 The Flat World and the Outsourcing Era

The internet and telecommunications revolution "flattened" the world, enabling knowledge work to be done anywhere on the planet. This gave rise to the Business Process Outsourcing (BPO) industry – a legacy approach to getting work done more cheaply by tapping labor in lower-cost regions. Companies in North America, Europe, and other high-wage areas discovered they could outsource routine processes (data entry, customer service, or payroll processing) to service providers in countries where wages were a fraction of domestic rates. For example, in the early 2020s, an IT professional in Silicon Valley might earn $140,000 a year, whereas a similarly skilled worker in a lower-cost region earned a small fraction of that. BPO firms built their model on this wage arbitrage: they'd charge clients a rate in between, pay the offshore workers much less, and pocket the difference as profit.

This global labor arbitrage lowered costs, but it didn't fundamentally change how work was done – it simply changed where it was done. The same process definitions and repetitive tasks were now executed by a person in Bangalore or Manila instead of Boston. Organizations did gain efficiency, but the outsourcing model often came with new friction: time zone differences, communication overhead, cultural and language barriers, and layers of extra management. Moreover, from a labor perspective, outsourcing often undervalues the workers actually doing the work. For every dollar a client pays to a BPO provider, 15–20% is profit and only a small fraction goes to the frontline workers. Workers were paid modestly for their time, while the value they produced was captured elsewhere. In essence, outsourcing stretched the traditional model of paying for time – it made labor cheaper, but still fundamentally billed by the hour or task, not by the outcome.

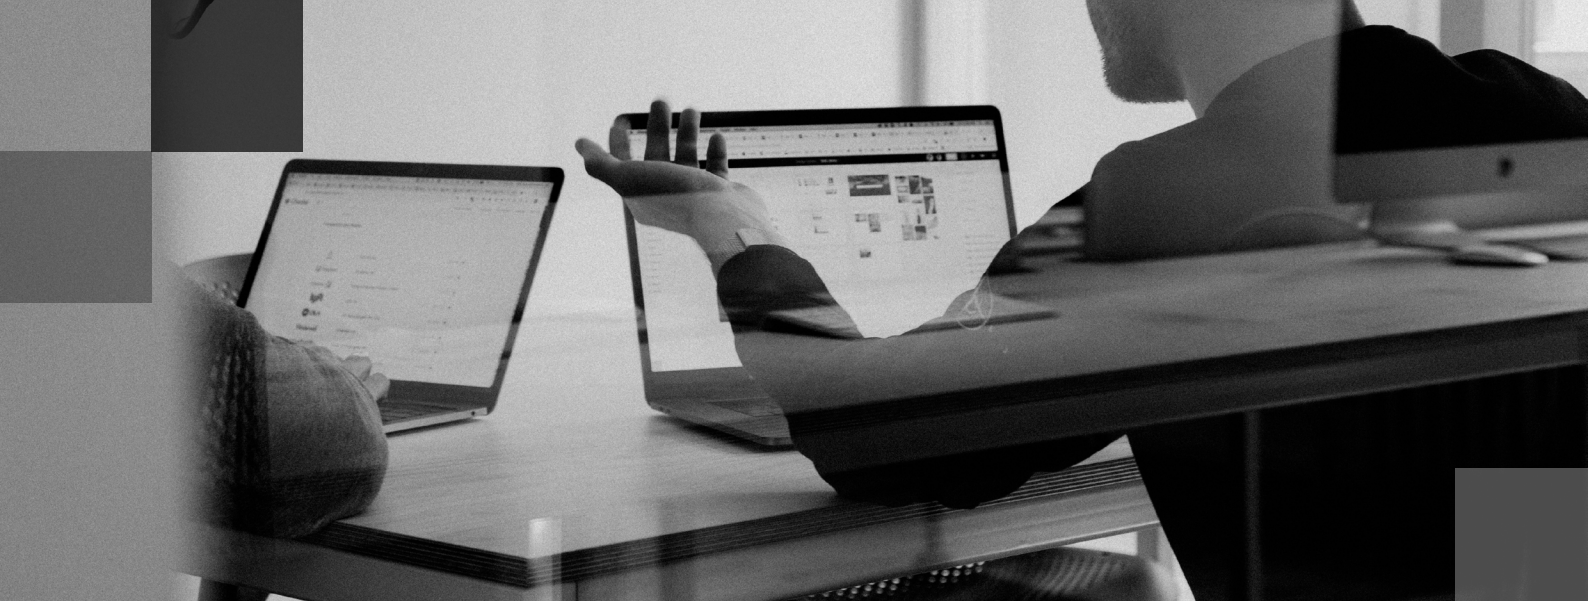

However, the outsourcing era also did something valuable.

It forced companies to explicitly define their processes and workflow endpoints.

To successfully hand off a piece of work to an external firm, businesses had to identify the inputs and outputs of each process, establish clear procedures, and set Service Level Agreements (SLAs) for quality and turnaround. In doing so, organizations created playbooks and documentation for countless workflows. On top of being an efficiency strategy of the old paradigm, the BPOs also laid the groundwork for the next evolution. By standardizing and codifying processes, companies made it easier for those processes to be taken over by automation.

By the 2020s, cracks were showing in the BPO model. Wage inflation in traditional outsourcing hotspots, growing concerns about data security and compliance, and the COVID-19-driven shift to remote work all began to erode the advantages of offshoring. If everyone is working remotely anyway, the distinction between a team in Wichita versus a team in Wipro begins to blur. Companies started asking: instead of sending work around the world chasing lower wages, could we use technology to collapse the cost of work even further?

In hindsight, BPO was a transitional phase. It stretched the old labor-cost paradigm to its limit while preparing the world for a more radical shift. The world learned to treat processes as modular, exportable units of work. By documenting workflows and proving that processes could be disaggregated and handed off, the outsourcing era paved the way for automation. The next step would be to automate those units.

1.3 Adding Intelligence to Automation

Automation in business processes is not new. For decades, we have used software to streamline workflows: from macros and scripts to full-fledged enterprise systems. In recent years, tools like Robotic Process Automation (RPA) have acted as "digital workers" following scripts to handle repetitive on-screen tasks, and enterprise software (ERP, CRM, etc.) has digitized many operations. These solutions delivered efficiency gains by removing some manual steps, but they had limits. For a start, **they are typically deterministic** – it does only what it's explicitly programmed to do and struggles with ambiguity or exceptions. Anyone who's interacted with an automated phone menu can relate to how brittle purely automated solutions can be when something unexpected happens.

The missing piece was **intelligence**. Until recently, if a task required judgment – even something as simple as recognizing a pattern in an invoice or understanding a free-form email request – automation often failed. Thus, humans remained in the loop for all these "edge cases" or for tasks too complex to distill into hard-coded rules. The next frontier of productivity is to automate not just the easy, structured tasks, but also the intelligent processing and decision-making parts of workflows, while still preserving human oversight where it truly matters.

Artificial Intelligence is now providing that missing piece. Machine learning and large language models are now handling a remarkable range of unstructured cognitive tasks: reading and interpreting documents, drafting responses, making predictions, and even conversing with customers. However, the goal is not an AI free-for-all. The goal is Supervised Automation – automation with a human supervisor in the loop to ensure quality and handle the hard parts. It's like an airplane with an autopilot: the autopilot handles the routine flying, but a human pilot is always there to monitor and take control in tricky conditions. In work processes, AI can now serve as that autopilot, and humans become the skilled pilots overseeing the journey.

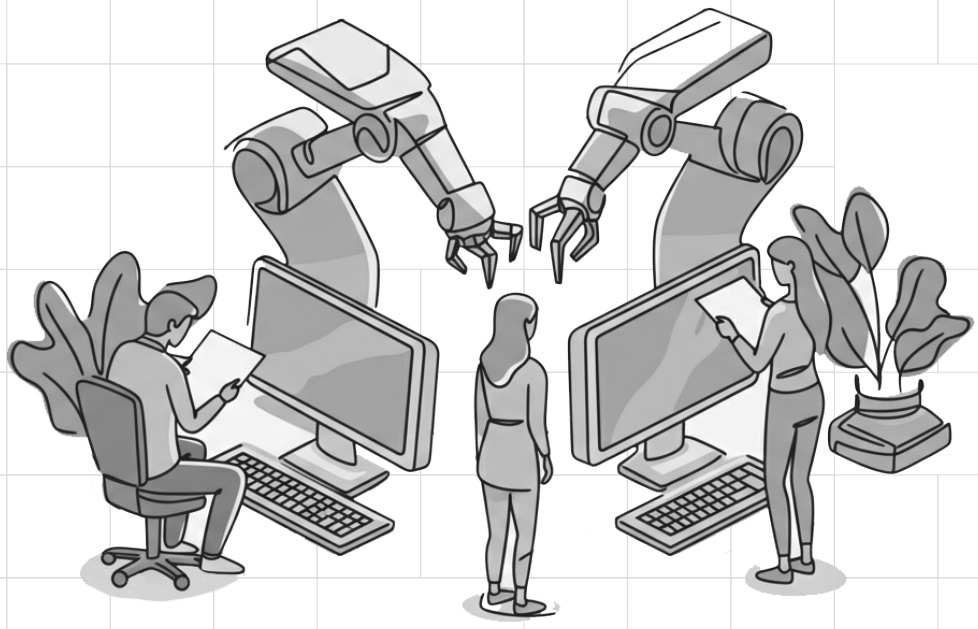

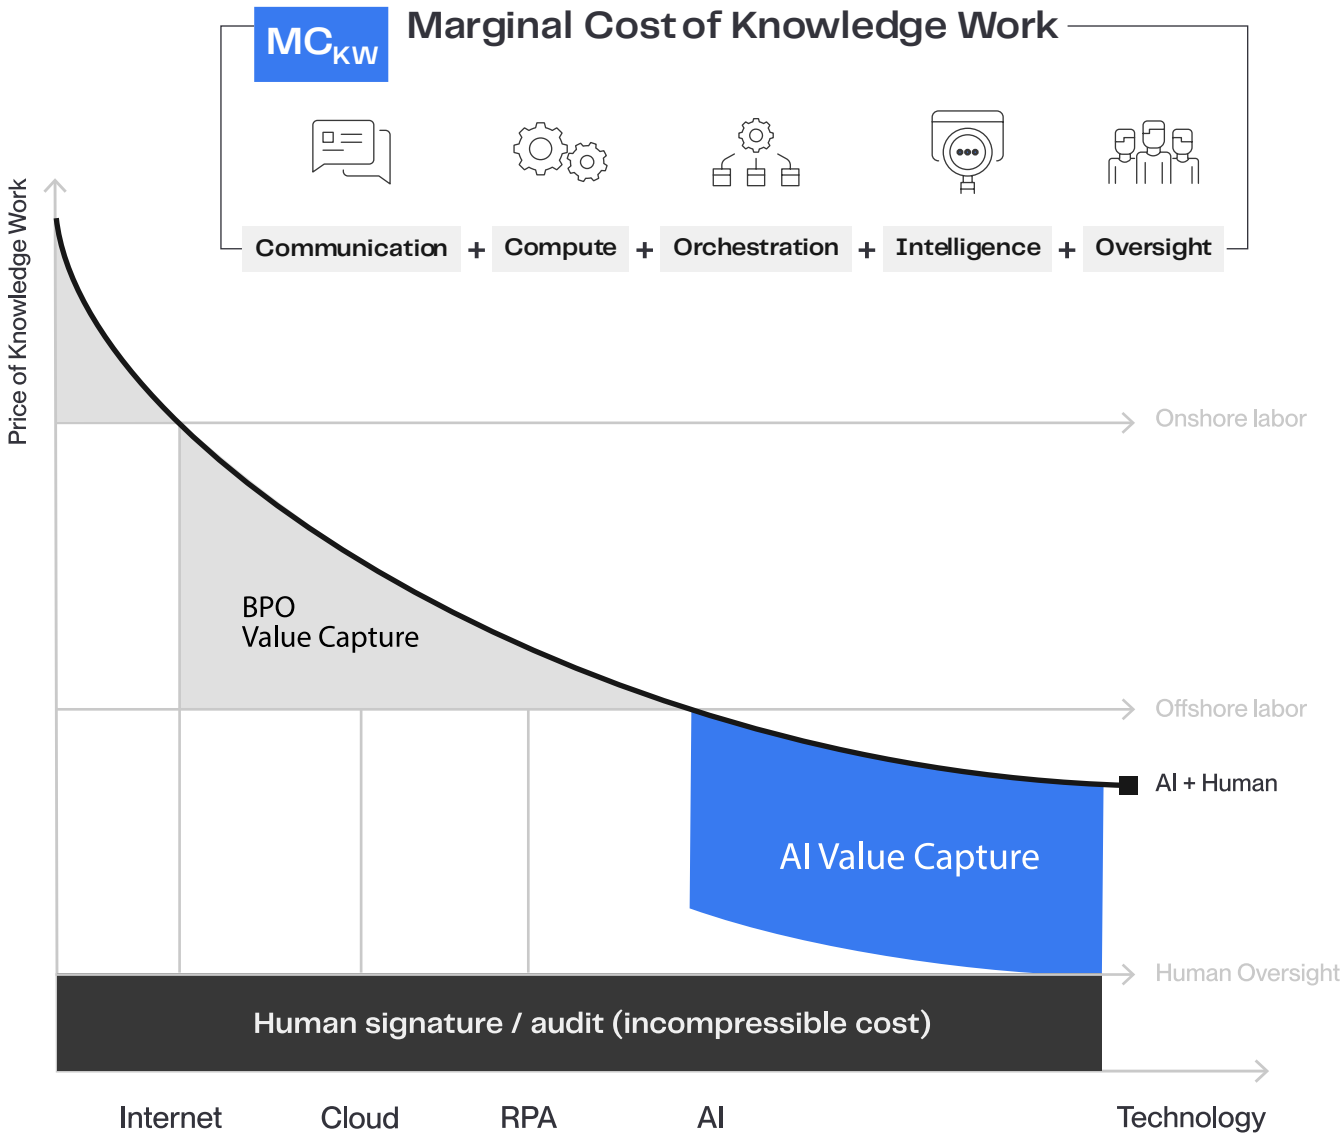

AI-driven automation is dramatically lowering the **marginal cost of knowledge work.** Where an offshore junior analyst might cost $3 per hour and a domestic one $15+ per hour, an AI performing a similar task can do so at an effective cost that is even lower – potentially mere cents per hour of compute– without sacrificing speed or quality. It collapses the cost structure of many business processes and upends the established geography-based wage advantages.

Cost is only one side of the story. The other side is quality and flexibility. The ideal supervised automation system operates 24/7, scales on demand, and maintains consistent quality at high speed while also incorporating humans at the right points. The advent of AI capable of cognitive tasks simply changes the roles people play and the scale at which they can contribute. As AI takes over routine work, we are witnessing a structural drop in the cost of knowledge work and a rise in productivity. As Sam Altman noted, the "cost of intelligence" is being driven to near-zero.

We stand on the brink of an era where labor productivity increases in orders of magnitude, as AI enables each human worker to be 10× more productive in terms of output, quality, or both. The stage is set for a transformation in how we work and how we value work.

Supervised Automation in Reality

In traditional thinking, automation was intended to replace human labor: if a task is automated, a person is removed, and vice versa. Supervised Automation turns this on its head. It posits that the best results come from a partnership – AI handling the repetitive heavy lifting, and humans providing guidance, expertise, and judgment at critical junctures. This is a bold promise: vastly higher productivity and better job satisfaction for workers are achieved simultaneously. How is that possible?

2.1 Humans + AI: Reimagining the Division of Labor

Start with a simple analogy: think of the work in your job as a meal. The creative parts – strategic planning, innovation, personal connections – are like the protein. The structured, routine parts – filing forms, updating data, managing processes – are like the carbs. Both are essential in a balanced diet, providing different but complementary benefits. In today's offices, highly skilled people often spend a significant portion of their time on the "carbs" of work, which are necessary but may not always leverage their unique talents. What if we could create a partnership where machines handle more of the carbs, freeing people to focus more on the protein? Supervised Automation creates this balance, not by eliminating carbs, but by redistributing where they're processed. As this transition evolves over time, both elements remain important – we're simply changing how the meal is prepared.

In a Supervised Automation model, AI systems handle the bulk of repetitive, structured tasks within a workflow. They digest documents, fill out forms, generate first-draft analyses, and route information – all the steps that don't require deep human insight. These AI agents work quickly and tirelessly. But unlike in a fully automated pipeline, humans are injected into the loop at key points to ensure everything stays on track. Think of these humans as expert reviewers or "judges." Their job is not to do the grunt work, but to verify and guide AI. Whenever the AI is uncertain, it flags the job for human review; whenever a decision is subjective or has legal/ethical implications, a human makes the call.

The result is a workflow that is mostly automated in execution, but supervised and audited by people. It's as if every worker has a team of tireless robotic assistants, and their new role is to lead and coach this team of bots.

This model dramatically increases a single human's leverage. If 95% of a process can be automated, one person can oversee that process across 20× the volume of work they could handle by themselves. To put it another way, the percentage of human work in each individual task may drop to, say, 5%, but that human's capacity scales across many more tasks or workflows – potentially one person doing the work of twenty by value, with AI doing the rest. We've seen this in practice: for example, in an insurance back-office scenario, instead of ten agents each handling 100 claims, you might have one expert overseeing 1,000 claims processed by AI, stepping in only for the few exceptions. The net effect is a 20× productivity boost per person. This isn't theory – it's happening in early adopters of these technologies.

2.2 Mid–Mile Corrections vs. Last–Mile Failures

One might think you just let the AI do everything and have a person check the final output – a last-mile human touch. But this approach can be flawed, especially if errors earlier in the process compound into a big mistake at the end. Instead, Supervised Automation emphasizes mid-mile human-AI "dances": integrating human checkpoints throughout the process, not just at the end. By catching uncertainties or subjective calls in the middle, errors can be corrected before they propagate downstream.

These mid-process expert reviews act as circuit breakers to catch errors before they compound. This is the "trust but verify" design principle: trust the AI to do the heavy lifting but verify at key junctures. If a mistake slips through one step, it gets caught at the next checkpoint rather than polluting the entire workflow. By catching mistakes early, the final output delivered to the client is reliable. **In this way, mid-mile corrections prevent last-mile failures.**

From a quality and risk perspective, this approach is transformative. It means AI can be deployed in high-stakes processes with confidence because there are always humans guarding against the edge cases or the unpredictable. This is how you avoid the nightmare scenario of a single AI glitch causing a massive, costly error. Supervised Automation is about trusting AI to work autonomously on 90–99% of tasks, but verifying its work whenever there is a doubt.

The result is not only faster work but better work – higher accuracy than either humans or machines could achieve alone.

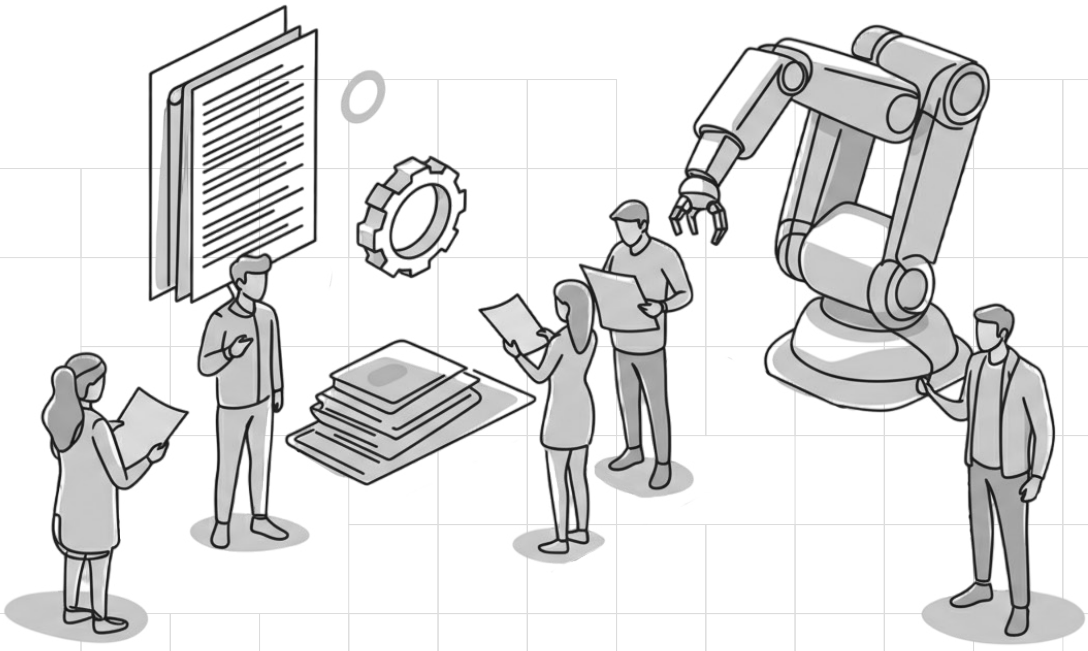

2.2 Outcome–Based Workflows and the End of “Hourly” Thinking

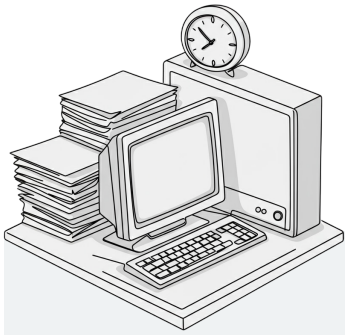

When AI and humans collaborate in this way, something profound happens to how we value work. We can stop counting hours and start counting outcomes. In a supervised automated workflow, it doesn’t matter whether a task took 2 seconds of AI processing or 2 minutes of human review – what matters is that the task was completed correctly. This enables a shift to outcome-based pricing and rewards. Instead of paying a BPO vendor or employee $X per hour, companies can start measuring cost per transaction, per document processed, and per case resolved. And because the **AI+human** system is so much more efficient, per-outcome costs drop dramatically while throughput skyrockets.

This outcome-centric mindset addresses a long-standing inefficiency: the fact that so many knowledge workers have been paid for attendance rather than results. It’s an uncomfortable truth that in many offices, output is only loosely correlated with hours at the desk. Some people produce brilliant work in a short time, others take much longer, yet both might be paid the same 8 hours of pay for a day’s work. By reuniting output with value, AI forces a more rational accounting. If one augmented worker can accomplish what used to take 10 workers, then that one augmented a worker becomes far more valuable to the company. We may even see compensation models evolve, where workers share directly in the productivity gains.

What’s important is that time stops being the sole metric of value. As AI resets the productivity equation, organizations that adopt supervised automation can achieve step-function improvements – not the typical 5–10% annual efficiency tweak, but 2×, 5×, or even 10× leaps in productivity.

In economic terms, it’s as if we suddenly moved to a higher production function altogether – a one-time jump thanks to a general-purpose technology, much like the jumps seen with electrification or the original IT revolution. But this time, it could be even bigger.

Solow Model: Productivity Jump with AI

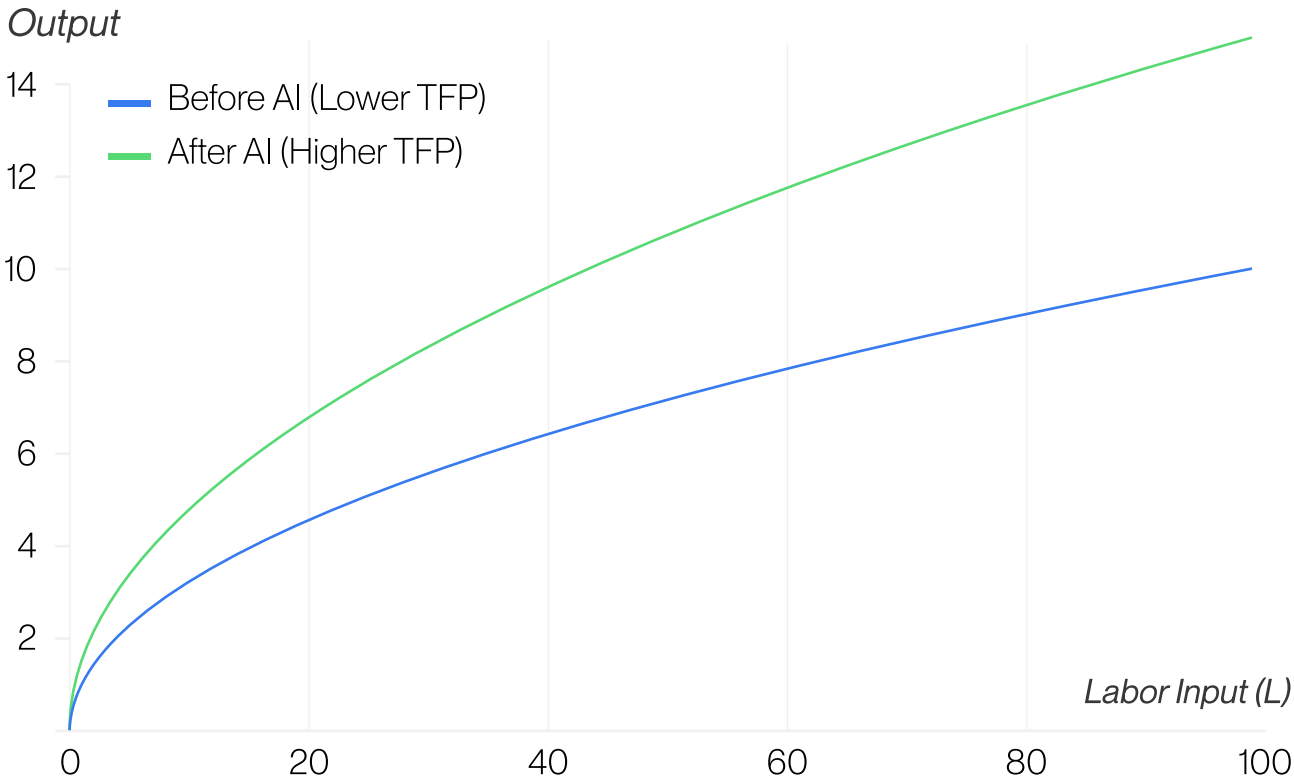

Solow Model: Productivity Jump with AI. Illustrative economic production functions show how a technological leap can raise output across all levels of labor input. In the graph above, the green curve represents an economy after a major AI-driven productivity boost, shifting significantly above the blue pre-AI curve. Such an upward shift reflects a higher total factor productivity (TFP): each unit of labor now yields more output than before. This is the kind of transformation supervised automation promises – a once-in-a-generation jump in the efficiency of converting human effort into results.*

We stand at the cusp of what could be a Solow-like boost to economic output: after years of slow growth and paradoxes, AI might finally be the technology that shows up emphatically in the productivity statistics, shifting the curve upward instead of just inching along.

The promise of Supervised Automation, therefore, is not just doing things faster or cheaper; it's redefining what we choose to do at all. When every routine task carries essentially negligible cost, businesses and individuals can refocus resources on innovation, customer experience, and creative endeavors. We move from a mindset of scarcity ("there's only so much labor, and we must allocate it carefully") to a mindset of abundance ("basic work is so cheap it's almost free – what ambitious projects will we tackle next?"). This can unleash a wave of experimentation. Projects that were once shelved because they were too labor-intensive or expensive can be revisited. Lean teams can afford a level of operational support that previously only large organizations could. A government agency can clear backlogs that citizens thought would take years.

In short, when the cost of the "carbs" of work plummets, we are freed to pursue more of the "protein" – the truly enriching, value-creating work.

opus

Human+AI in Practice

3.1 Workflows as Dynamic, Learning Systems

The central architecture of Opus is the idea of a workflow – a series of well-defined steps that transform inputs into outputs. Every workflow is represented as a living, breathing digital entity, specifically, as a Directed Acyclic Graph (DAG) of tasks. Each node in the graph is a modular skill or task (like “extract data from form” or “send approval email”), and edges dictate the sequence or dependencies. Some of these tasks are fully automated, and some are human expert review tasks where a person must make a decision or verification.

This structured representation is powerful: it makes the process transparent and auditable since every step and decision point is explicitly laid out and logged. It also makes the workflow modular – you can improve or swap out one component without needing to redesign the whole process.

Opus’ core intellectual property includes the Work Knowledge Graph (WKG) and the Large Work Model (LWM) – advanced technologies that enable the creation and optimization of such workflows. The Work Knowledge Graph is essentially a vast network (graph) capturing procedural knowledge from many industries – a kind of encyclopedia of how various tasks can be performed and linked. It encodes best practices, regulatory requirements, and domain-specific steps gleaned from studying thousands of real business processes.

The Large Work Model, on the other hand, is a generative AI model that can produce workflows by drawing on the knowledge in that graph. When a user of Opus wants to automate a new process, they don't start from scratch; they provide an Intention – essentially a high-level specification of what they want to achieve (the input, output, and context for the process). Given this intention, the LWM queries the Work Knowledge Graph and generates a tailored workflow that accomplishes the goal.

This approach allows rapid deployment and iteration of automated workflows. It's not a months-long software project; it can often be done in minutes or hours on the platform because the heavy lifting of understanding process logic is aided by AI.

Once a workflow is up and running, it becomes a learning system. Just like Google Maps improves over time by learning from traffic patterns and user corrections, an Opus workflow improves by learning from each execution and outcome. Every time a human expert intervenes, that scenario is recorded. Engineers and the AI itself can then update the logic or model to handle that case automatically next time. Over hundreds or thousands of iterations, the number of edge cases that require human help tends to shrink – the workflow "learns" to handle more on its own. Workflows that start at 70% automation rate can climb to 90–95% automation after a few months of feedback and tuning.

This is fundamentally different from traditional outsourcing or even earlier automation, which remained static until a human redesigned them.

How does Opus deal with sensitive data or industries with strict data sovereignty? The platform can be run in cloud or on-premises environments as needed, and specific AI models can be swapped to meet requirements. Since every step can be permissioned and audited, organizations can ensure that only certain vetted humans access data, or that all data stays within a jurisdiction. All of this happens under governance. **Every step and decision in Opus is recorded.** There's a complete audit trail for each workflow run, which not only helps in oversight but also in continuous improvement. This full auditability and traceability is crucial for regulatory alignment. Highly regulated industries can adopt supervised automation with confidence because nothing is a black box.

You can always answer who did what, when, and why.

3.2 The New Role of the Human: Expert Reviewers and AI Trainers

In a fully realized supervised automation environment, what do humans do all day? This question naturally arises, especially with the specter of "automation = job loss" looming in many minds. The answer is optimistic: humans do more important work.

First, within the workflows themselves, humans act as Expert Reviewers at the points requiring judgment, empathy, or approval. We deliberately use the term "expert" because this is an elevated role – it's not a mechanical task to mindlessly review what the AI did; it's a skilled task of making decisions on nuanced cases. The people performing these reviews are empowered to use their domain knowledge and experience. In a legal contract review, AI can mark up standard clauses and highlight anomalies, but a lawyer reviews the few non-standard elements and decides if they're acceptable. These human decisions are fed back as learning. Humans become the teachers and supervisors of fleets of AI workers – AI trainers and AI overseers.

Outside of direct workflow execution, new human roles emerge as process designers, AI orchestrators, and exception handlers. Business teams that used to draft lengthy requirements for IT can now directly craft and tweak their workflows in platforms like Opus through intuitive interfaces – a bit like building a flowchart with smart suggestions. This empowers the actual process owners to implement improvements on the fly, under guardrails set by IT for security and integration. It dramatically reduces the cycle time of process re-engineering from what might have been a 6-month consulting project to an afternoon of dragging and dropping components. The high capital expenditure of old-school business process reengineering – hiring consultants, writing custom code, overhauling systems – is collapsed into a software service where the heavy lifting is done by the platform.

This not only speeds up innovation, it also leads to better outcomes: since the domain experts are directly at the controls (with AI copilots and IT safety nets), the resulting workflows tend to be more accurate and aligned to the business needs.

This reunites decision-making with execution. The people who understand the work are now the ones empowered to configure how it gets done by the AI.

Moreover, a new kind of workforce is emerging: a human expert cloud accessible on demand. If your automated workflow needs a human review at 3 AM, you might not have someone in your office at that hour – but somewhere in the world, an expert is awake and available. Platforms like Opus can route those review tasks to a distributed network of professionals, almost like how Uber routes ride requests to drivers. But unlike Uber, which is constrained by physical cars and local presence (the realm of atoms), an expert cloud operates in the realm of electrons – instant digital connectivity. It's not limited by geography; an expert in Australia can review a task for a U.S. company while the U.S. team sleeps. This makes the expert layer extremely liquid and scalable. During surges of work, more experts can be pulled in; during lulls, they can be doing other things.

It's a far cry from the rigid staffing of traditional operations. Some have compared this model to an "Uber for knowledge work," but in truth, it's even more powerful because of the lack of physical constraints. We're essentially virtualizing expertise. Professionals can be plugged into automated workflows anywhere, anytime.

Crucially, this human expert cloud is global and meritocratic. Work can be routed to whoever is best qualified and available, regardless of location or background. It doesn't matter if the expert is in New York or Nairobi – if they have the skill, they can contribute. This opens up opportunities to people everywhere and creates a truly merit-based talent layer on top of the world's workflows. A complex claim review might go to a former insurance adjuster working from home in Manila; a medical transcription check might go to a nurse in Mumbai. All experts are connected through Opus, and all are judged by the quality of their work, not by their location or proximity. In effect, supervised automation platforms work for any capable worker in the world.

For businesses, this on-demand expert layer provides unprecedented agility. If you suddenly have double the workload, you don't need to hurriedly hire and train new staff or request your team for overtime; the network of experts can scale up instantly to absorb the spike. If volume drops, you're not stuck carrying excess staff – the experts simply get routed other work elsewhere. It's a highly efficient allocation of human talent. From the worker's perspective, it provides new avenues to earn and participate in the global economy.

A skilled individual can contribute to multiple organizations' processes in a single day, choosing tasks that fit their expertise and schedule.

How will these changes affect everyday work? Perhaps the most important shift is that employees in an enterprise using supervised automation become managers of their processes. They are not just cogs performing tasks; they are overseers and improvers of tasks performed by AI. They can continuously ask, "How can this process be better?" and implement changes quickly. This turns operations teams into innovation teams. The work culture shifts from one of executing routine procedures to one of managing outcomes and refining processes. It's empowering and motivating: people feel a greater sense of ownership because they're responsible for the outcome, not just a narrow task. They also gain transferable skills (like orchestrating AI workflows) that will be highly valuable in the future.

opus

A New Workflow Engine ■

How do we bring about this future of work?

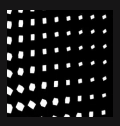

At its core, Opus is a workflow automation platform with AI as its engine. When a process owner wants to automate a workflow, they can simply start by describing what they want to achieve. The intention of the workflow – essentially the input, process, and desired output – is then captured in a structured way by Opus.

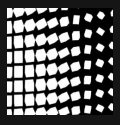

The Opus Large Work Model interprets this intention and searches within the Work Knowledge Graph for relevant process knowledge. Within seconds, Opus drafts a new workflow: a series of tasks, both automated and human, that should achieve the goal. The user can then refine this if needed, perhaps adding company-specific rules or integrating specific software tools like emails, databases, and enterprise apps.

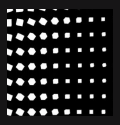

Once the workflow is configured, Opus deploys it. The tasks designated for AI are executed by Opus' own AI models or by calling external AI services as appropriate. The tasks designated for human expertise are routed either to the company's own employees or to Opus' on-demand expert cloud, depending on the process owner's preferences. Every step flows through the Opus platform, which orchestrates the entire process and collects data on it.

One of the transformative aspects of Opus is how it democratizes access to process automation. You do not need to be a software engineer to create a complex workflow. The platform provides a visual interface and AI guidance, so business users can drag-and-drop to modify sequences, set rules, and deploy changes instantly. This is all done under the governance of the IT department's guardrails.

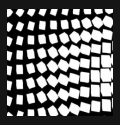

By handling these concerns centrally, Opus ensures that compliance and security are baked in. This is essential, especially in highly regulated industries. Every input, output, and intermediate decision is recorded in a complete audit trail, which not only aids oversight but also supports continuous improvement.

3.3 Manhour as the Atomic Unit of Work

Opus introduces a novel approach to valuing knowledge work automation, utilizing the manhour as its fundamental unit as a basis for value-driven pricing and applying this familiar labor metric as a standardised unit of AI-driven productivity. This unit – much like "horsepower" was coined to relate engine power to a familiar concept (horses) or "wattage" for electrical output – represents a normalized measure of work completed by the Opus system, whether through AI, human expertise, or a blend of both, serving as a clear proxy for the value generated.

This pricing model flips the typical software sales approach. Rather than charging per software license or via a subscription fee for a tool that then requires you to put in the effort, Opus charges for work done. By establishing a clear unit of work output (the manhour), Opus provides a simple benchmark for organizations. The value proposition becomes immediately apparent when contrasted against traditional labor arbitrage models – whose own cost per manhour often exceeds Opus' price – or fully loaded onshore costs.

What does an Opus "manhour" include? It includes all the AI processing and all the behind-the-scenes orchestration. Essentially, Opus takes on the complexity of combining AI and human effort and gives the client a single metric. This value-driven pricing, measured via manhour output, aligns incentives strongly: the client pays only for actual productivity, and Opus is incentivized to complete work efficiently and accurately. It's an exponential simplification of cost calculus for the client – providing clear ROI visibility as you know exactly what each unit of work will cost, and it's a fraction of traditional methods.

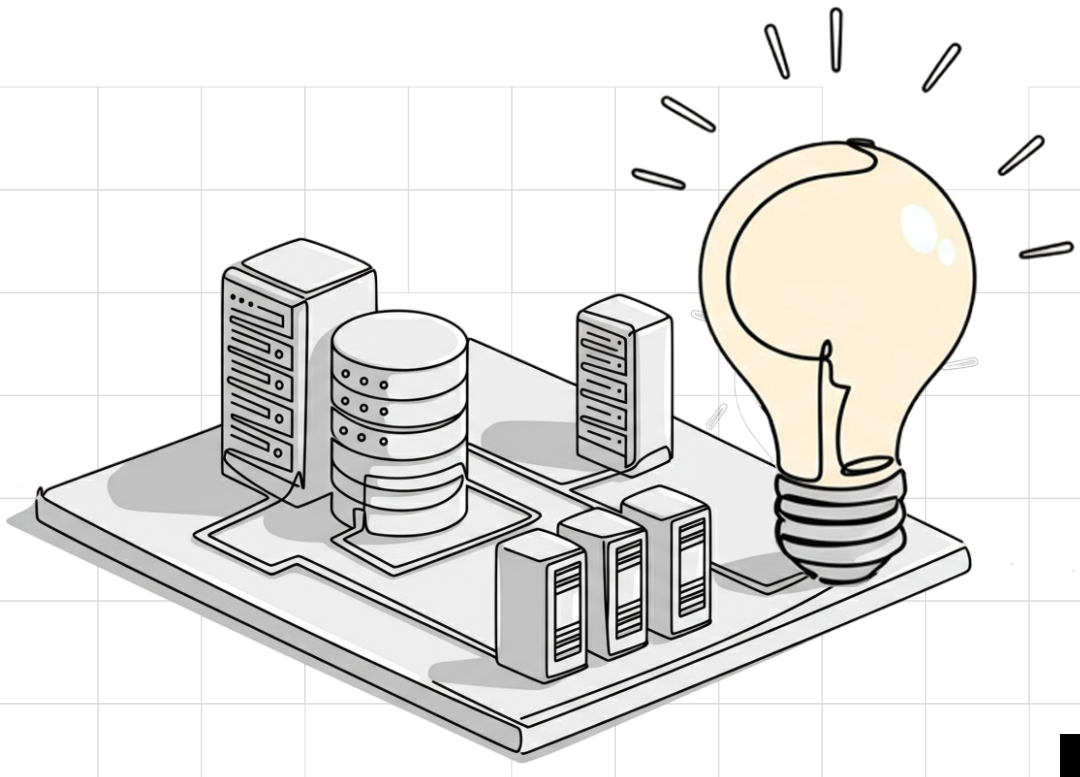

3.4 Reuniting People with Purpose

Perhaps the most important thing Opus does is less tangible: it helps reorient organizations to **think in terms of outcomes and purpose** rather than tasks and toil. By trusting AI to handle the grunt work, companies send a message to their employees – we value your time, and we want you to use it on what really drives the business. It's inspiring for a team member to know they implemented automation that saves their colleagues hundreds of hours, freeing those colleagues to do more meaningful projects.

In numerous deployments, we've seen initial skepticism ("Will this AI take my job?") turn into enthusiasm ("This AI makes my job so much better!") once people experience the shift. Workers become more like product managers of their processes, constantly thinking of improvements and new features (since they can implement them via Opus), rather than feeling stuck executing the same steps day after day.

Opus effectively collapses the cost and barrier of innovation within operations. It's not just automation, it's rapid experimentation. A team can try a new workflow or a tweak, see the metrics (Opus provides real-time dashboards of throughput, error rates, human involvement rates, etc.), and either roll it back or roll it out further. This agility in operations was unheard of in the past when any process change meant lengthy retraining or new software development. Now it might be as simple as a version update in a workflow diagram. This means operations teams can iterate and improve continuously, just like software teams have done with Agile methodologies.

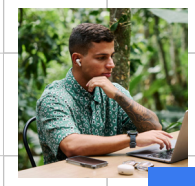


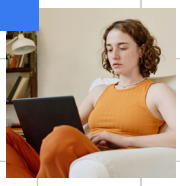
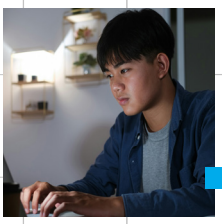

Finally, Opus hints at a future labor model: a marketplace of knowledge work. As more companies use platforms like this, the boundary of the firm starts to blur. Why couldn't a skilled individual in Spain be an "expert reviewer" for a U.S. company's Opus workflow as part of their day, essentially plugging their expertise into multiple organizations on their own schedule? This is already happening on freelance platforms, but with Opus, it can be more seamless – tasks find the expert rather than experts bidding for tasks.

We foresee a world where anyone with talent can opt-in to contribute to workflows globally, and companies can source talent not just by hiring full-time employees or outsourcing to BPOs, but by tapping into a liquid expertise marketplace enabled by AI workflow platforms.

This democratizes opportunity: a knowledgeable person in a small town can participate in high-value work for a big corporation without relocating, and the corporation gets just-in-time access to their skill without a permanent hire. Opus is building toward this vision carefully, making sure quality and trust can be maintained (through vetting of experts, encryption of data, etc.), but the technological foundation – the work graph, the AI orchestration, the auditability – is already in place to support it.

In summary, supervised automation doesn't eliminate the human – it elevates the human. It changes the question from "Will AI take my job?" to **"How can I expand my job with AI?"** The managers, analysts, and specialists of tomorrow will be those who know how to direct armies of AI helpers, how to integrate their expertise at key moments, and how to design workflows that harness the best of both AI and human capabilities.

An Abundant Future Anchored in Values ■

The trajectory laid out in this paper points to a remarkable future – one where AI ushers in a world of abundance in terms of productivity and possibilities. By collapsing the cost of routine work and amplifying human capabilities, Supervised Automation challenges the old assumptions of economics that we must always make trade-offs due to scarcity. When the "robotic" portion of work becomes essentially free and infinitely scalable, the only truly scarce resource remaining is human time and attention. This has profound implications. It means that how we choose to spend our time becomes more important than ever. It also means that societies and organizations need to rethink how value and wealth are distributed because the link between hours worked and value produced will be increasingly tenuous.

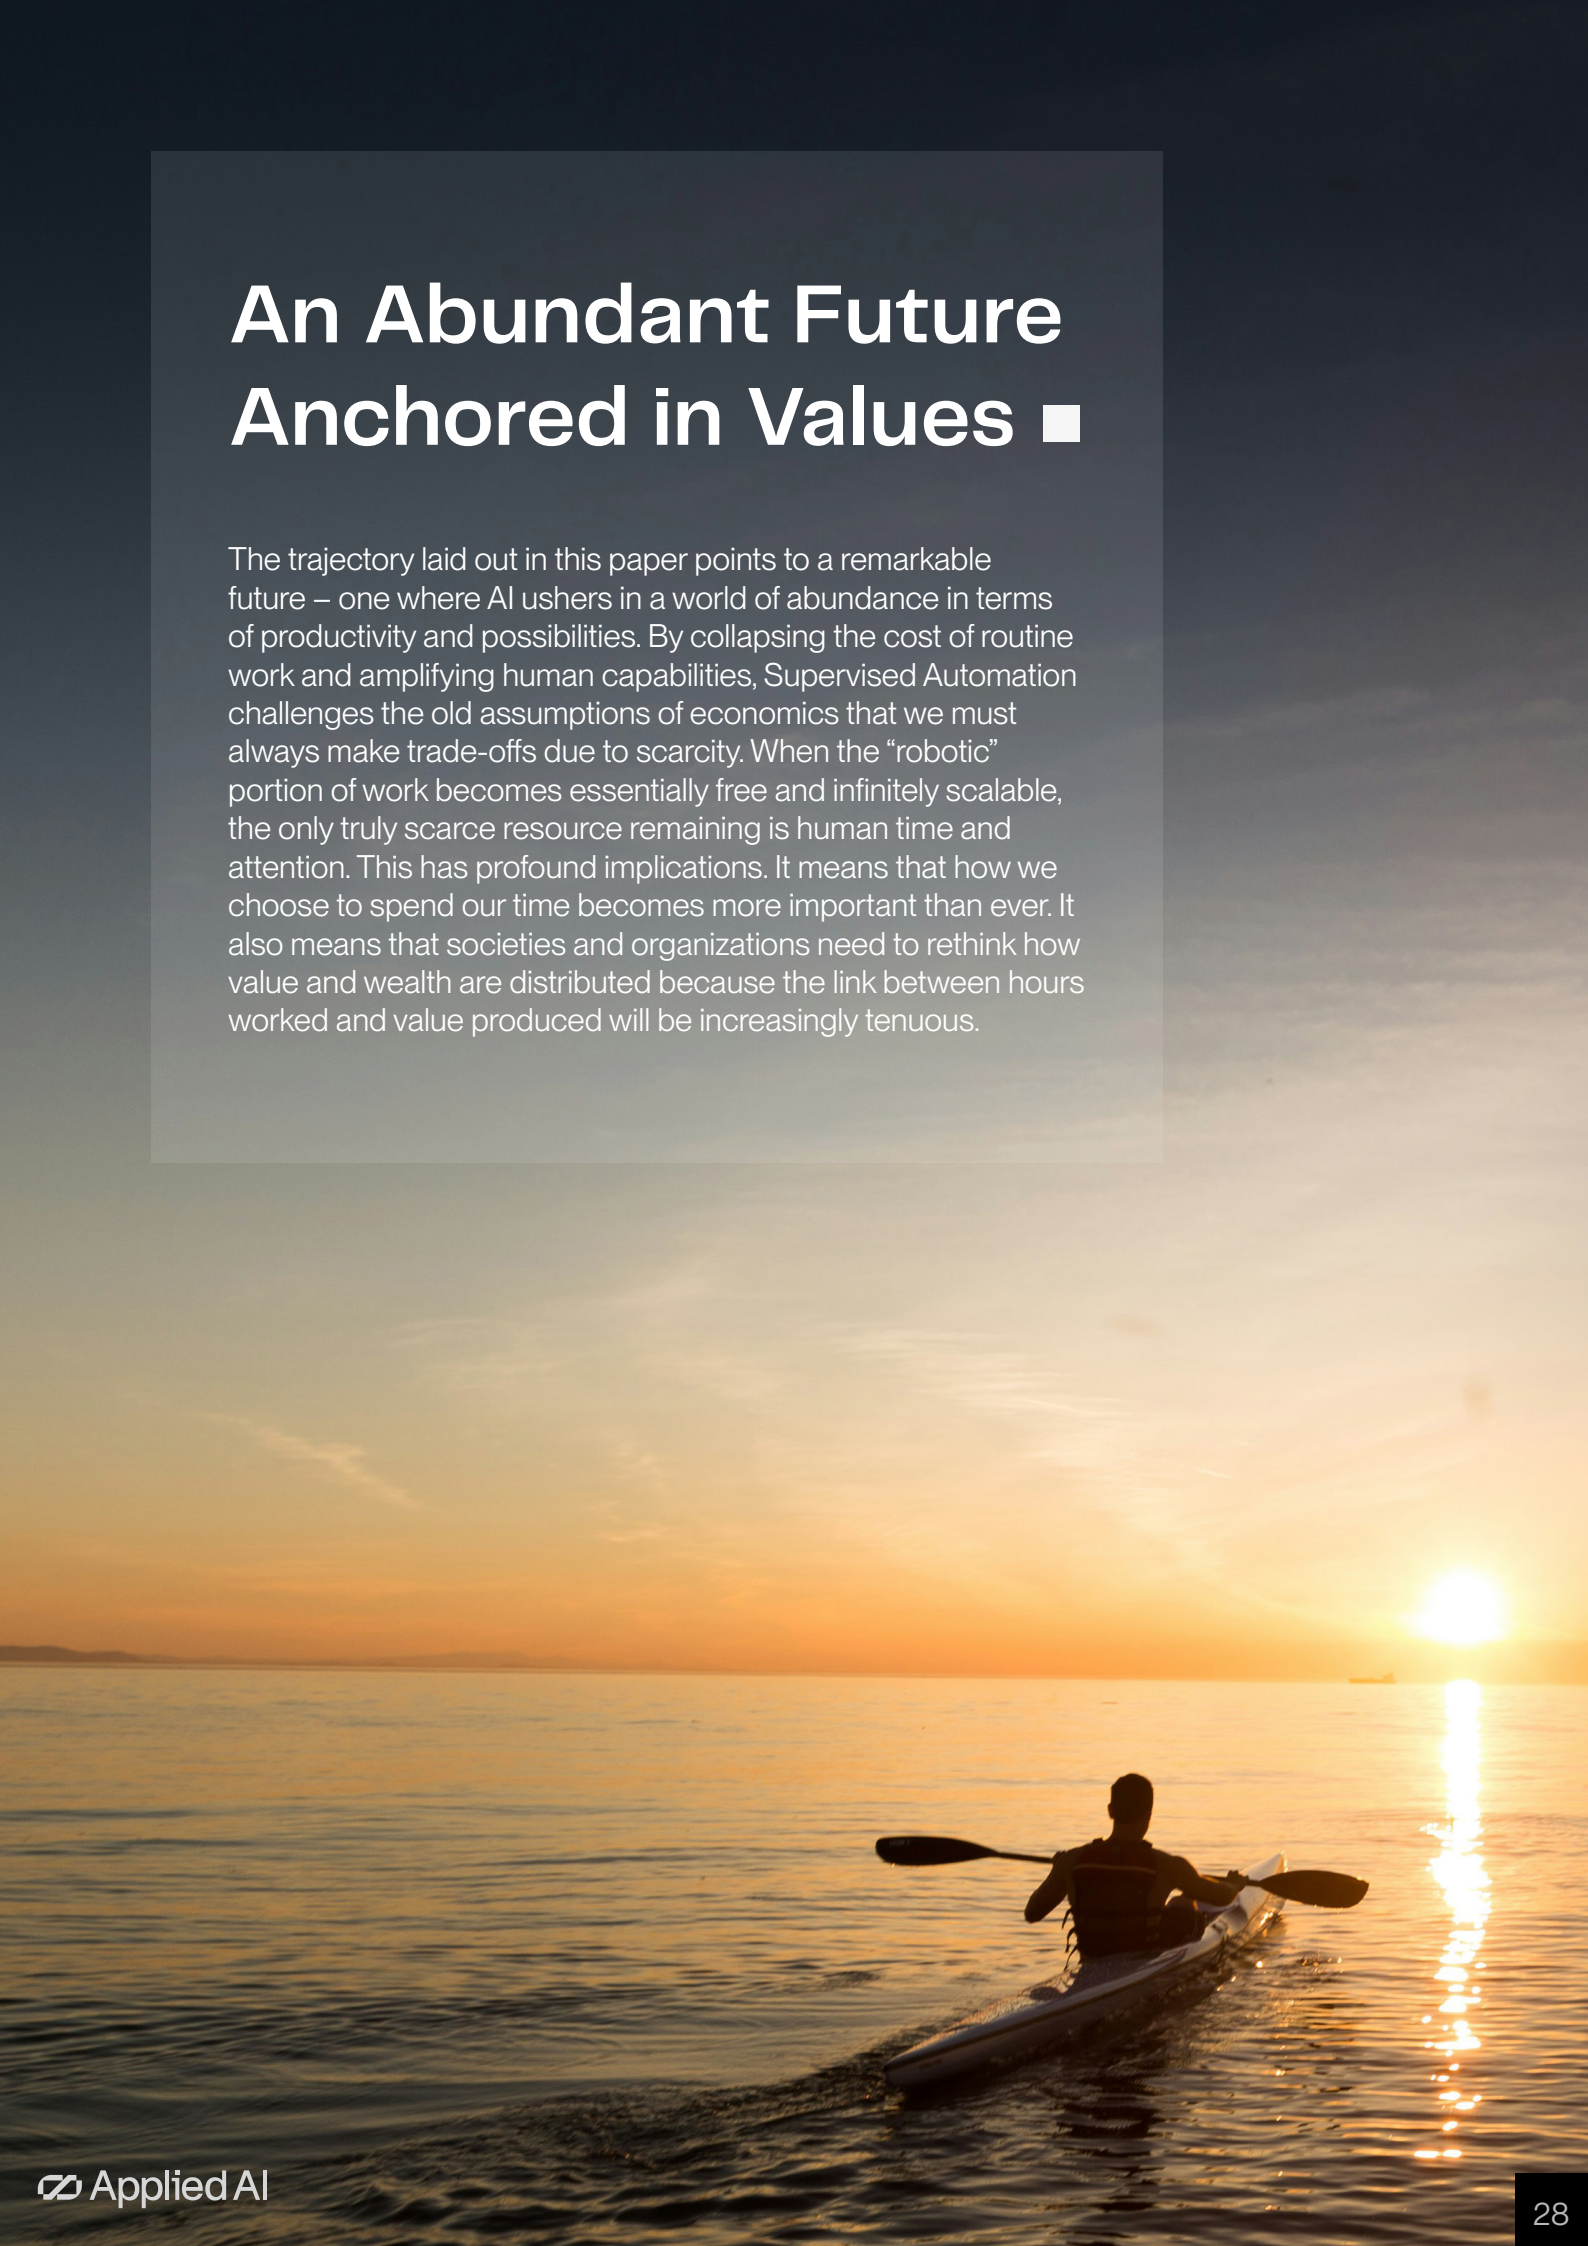

4.1 Beyond Scarcity: Rethinking Economics and Work

Classical economics is built on the premise of scarce resources – labor, capital, and raw materials. But if AI can provide effectively unlimited labor for many tasks, we move toward an "economics of abundance." The bottleneck shifts to human creativity, innovation, and oversight. We will never run out of problems to solve or new ideas to pursue. This isn't wishful thinking – it's the logical outcome if productivity increases several-fold across knowledge industries.

However, the transition period is crucial. Productivity gains need to translate into widely shared prosperity, not just concentrated profits. In the Industrial Revolution, productivity eventually lifted living standards broadly, but not without social turmoil, policy interventions, and the empowerment of workers. Marx's concept of surplus value – the idea that in a capitalist system, workers often create more value than they are paid for, with the surplus accruing to the owners of capital – is highly relevant when the "workers" increasingly include AI machines. Without deliberate action, all that value would go to the companies deploying AI, and very little to the human who is now supervising 20 workflows. This could exacerbate inequality and lead to a new form of worker alienation, where people feel even more disconnected from the fruits of production.

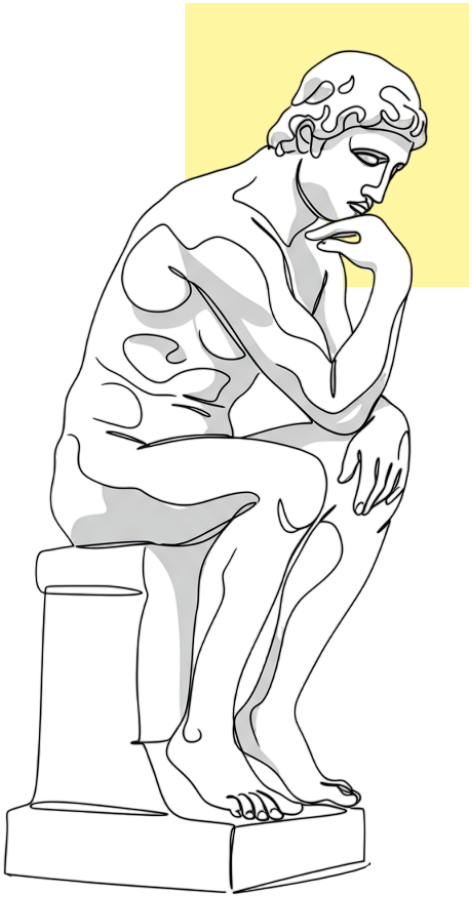

It doesn't have to be that way. We have the opportunity to design models where the benefits of automation are shared.

This could mean equity or profit-sharing for workers who oversee AI; it could mean new forms of collective ownership of AI tools; or policy mechanisms like taxes on AI productivity that fund social dividends. From a business perspective, forward-looking companies are already thinking about retraining and upskilling their workforce so that employees displaced from one role by AI can move to higher-value roles – often enabled by the same AI. In many of our deployments, not a single person lost their job; instead, their job changed, and often for the better. But scaling that outcome economy-wide will require conscious effort.

Economists have long noted that productivity isn't everything for society – distribution and use of productivity gains matter too. If AI enables a 10× leap in output, that creates a big pie of wealth. How that pie is divided will determine social outcomes. There is a moral dimension: those who contribute (even if it's guiding an AI that does most of the work) should share in the benefits. There's also a pragmatic dimension: broad-based prosperity fuels a virtuous cycle of growth, whereas extreme concentration of gains can lead to instability and backlash.

The moral and legal centrality of humans must remain a cornerstone. Ethically, we believe every automated decision ultimately needs a human accountability. Legally, someone must own liability when AI makes a mistake (as it inevitably will at times). That "someone" can't be a machine – it will be an organization or person. Recognizing this leads to the design principle that humans remain in the loop not only in operations but also in responsibility. Supervised automation inherently respects this: **A human is always overseeing, and thus a human can always be answerable.** This will likely be reinforced by regulations – for example, laws might require that any AI decision affecting a customer (credit approval, hiring decision, etc.) has a traceable human sign-off or review.

In macroeconomic terms, if AI really does bring a large one-time jump in productivity (a new Solow residual in growth accounting), we should prepare for its effects. GDP per capita could surge in those nations that adopt it, potentially widening the gap with those that don't. This puts a premium on countries and companies to not fall behind in adoption. But it also raises questions: if GDP grows but employment doesn't need to grow as much, how do we keep aggregate demand high? Some economists have started discussing ideas like a shorter workweek or universal basic income (UBI) in the context of AI-driven abundance. If society can produce plenty for all with much less human labor, maybe we can work less and find more meaning, supported by the wealth the AIs generate. These are radical shifts from today's norms, but they may become practical conversations in policy.

4.2 Embracing the Future – Ethically and Proactively

To truly reframe the global debate on labor and automation, we should reject the zero-sum narrative of "AI vs jobs" and replace it with "AI for humanity's growth." Supervised Automation exemplifies that middle path where AI and humans together do more than either could alone.

We reiterate a hopeful vision: AI-driven supervised automation can create a world of abundance – more goods and services produced with less drudgery, solving pressing global challenges by unleashing human creativity. But this world will only be more prosperous and more human if we intentionally design it that way. Productivity, as the Nobel Insight reminds us, is the key to long-term prosperity, but how those productivity gains are used is up to us. Will they widen gaps or bridge them? Will they diminish work or enrich it?

Our stance is clear: we see a future where the percentage of human effort in many jobs might drop to 5% or 10% of what it once was, yet that effort will be spread across many more endeavors. A person might supervise dozens of processes and projects, effectively contributing to far more output than ever before. In doing so, they own their work more fully – they guide it end-to-end with their AI partners – and they derive greater satisfaction from focusing on strategy, creativity, and human connection. Meanwhile, consumers and society reap the benefits of cheaper services, faster innovation, and solutions to problems previously considered too costly or complex.

Opus is one piece of this puzzle – a platform proving today that this model works, that it's not science fiction but a present reality. Our hope is that by sharing this whitepaper, we spark a broader conversation and inspire action. The future of work is not a spectator sport; it's something we are all tasked with creating. With the boldness to adopt new tools, the wisdom to maintain human values, and a commitment to shared progress, we can indeed reframe the debate. It's not about humans versus machines. It's about humans and machines, together, unlocking potential. Productivity will soar, but more importantly, people will thrive – and that is the ultimate promise of the future of work.

Education and Upskilling. Workers should be trained not only to use AI tools but to excel in the areas AI can't. Skills like critical thinking, interdisciplinary problem-solving, emotional intelligence, and domain-specific expertise become even more valuable. Organizations and governments need to invest heavily in education and retraining so that people are prepared for these elevated roles. As growth models suggest, human capital is as important as technology in driving progress. We should treat AI literacy and collaboration skills as fundamental parts of human capital development going forward.

Policies for Transition. There may be turbulence in specific sectors. Safety nets, transition assistance, and perhaps more radical ideas like universal basic income could be considered to ensure no one is left behind in the short run. In the long run, if productivity gains are enormous and broadly distributed, society can well afford to support those dislocated while they find new paths. The key is managing the transition humanely, as was eventually done in past technological shifts via mechanisms like social security, public education, and labor rights.

Worker Representation. Just as labor unions fought for fair wages and conditions in the industrial era, we may need new mechanisms for the representation of workers in decisions about AI adoption. If a company is going to implement technology that multiplies a worker's output twenty-fold, workers should have a voice in how that affects their role, pay, and work-life. Rather than resisting the tech, employee representatives and leaders can negotiate to share in the benefits – shorter workweeks for the same compensation, or bonus structures tied to the productivity gains. The aim should be to align everyone's incentives toward making the transition a win-win.

Ethical Frameworks and Governance. Each organization deploying AI, and society at large, should have clear ethical guidelines. Bias mitigation, privacy protection, and decision transparency should be core requirements for any AI workflow. On our part, Opus has built-in bias checks and that comprehensive audit trail helps in auditing for fairness and compliance. We also anticipate regulatory frameworks will evolve that define responsibilities. These ensure that efficiency doesn't come at the cost of justice or trust. Proactive self-regulation by industry can preempt heavier-handed regulation and build public confidence.

Harness the power of AI to deliver business results.

opus

- Transform strategy into automated workflows in days
- Reduce operational costs while improving compliance
- Integrate human oversight at critical decision points
- Deploy enterprise AI while building in-house expertise

The Opus Difference

Redesign

Rapid Workflow Design Intuitive no-code tools for swift visualisation, creation, and modification.

Digital Twin Simulation Model and test process changes in a risk-free virtual environment.

Integrate

Centralized Hub: Connect point solutions and enterprise systems for seamless data and workflow handoffs.

Team Empowerment: Equip process owners with the visibility and capabilities to manage their workflows effectively.

Supervise

Intelligent Task Automation Leverage AI to execute repetitive process steps efficiently and perform regulatory and business rule checks.

Human Oversight Designate critical checkpoints for human review, validation, and decision-making.

Evolve

Performance Analytics & Insights Monitor key process metrics and identify areas for targeted change, and performance tracking.

AI Optimization Suggestions Receive process improvements suggestions based on performance data, subject to human approval.

KYC Verification

Banking

Automates straight-through processing for high-confidence cases while routing exceptions for human review.

80% reduction in manual time spent

4.5x increase in productivity

11% improvement in accuracy

Independent Medical Examination

Insurance

Organizes complex medical records to produce summaries of the most important events with full audit trails.

88% reduction in review time

7.0x increase in productivity

16% increase in accuracy

Pharmacovigilence

Pharmaceuticals

Monitors medical literature, media, and reports in multiple languages to generate adverse event audit reports.

84% reduction in review time

6.4x increase in productivity

18% increase in accuracy

Conclusion

The age of supervised automation is upon us, and it brings immense promise. It offers a path to resolve the long-standing tension between productivity and humanity in work. We no longer must choose either efficiency or meaning – with the right approach, we can have both. To realize this vision, every stakeholder has a role to play:

Enterprises: Take initiative before competitive pressures mandate change. Start identifying workflows that can be transformed and empower your teams to become architects of these new processes. By adopting this paradigm early, you not only gain in efficiency and cost savings, but you also position your organization as a leader in innovation. Treat your employees as partners in this transformation – involve them, train them, and share the gains. The companies that thrive in the next decade will be those that reunite work around outcomes and elevate human talent to its highest and best use.

Policy Leaders and Educators: Lay the groundwork for a smooth transition. Update curricula and training programs to focus on complementary skills for the AI era – creativity, critical thinking, AI oversight, and interdisciplinary problem-solving. Craft policies that encourage innovation but also protect people – this may include incentives for companies that upskill workers, modernized labor laws that account for human-AI collaboration, and social safety nets (like retraining grants or portable benefits) for a more dynamic work landscape. Importantly, regulators should require transparency and accountability in AI systems without stifling progress.

Knowledge Workers and Professionals: This new paradigm is not about replacing you – it's about amplifying you. Rather than fear AI, learn to leverage it. Cultivate the uniquely human skills that automation can't replicate easily. Become the expert who trains and supervises AI in your domain. Be curious and open-minded about new tools like Opus; experiment with them to automate the parts of your job you dislike, and spend the freed-up time on work you love or new initiatives that can set you apart. This is your chance to multiply your impact.

You can be more productive, more creative, and find greater purpose in your work when the rote busywork is handled by your digital assistants.

The future of work in the age of supervised automation is bright: it's a future where work is reunited with purpose, where human time is valued as precious, and where innovation can flourish freed from the shackles of routine. We have the tools and vision to avoid the mistakes of past industrial revolutions. By actively designing systems where humans and AI collaborate – keeping humanity in the loop – we ensure technology serves us, not vice versa.

The next chapter of work will be written by those pioneers who dare to leverage these new tools to their fullest. Let us embrace this bold vision with both excitement and responsibility. The transformation has begun – and with it comes the chance to redefine work for generations to come, unlocking productivity hand-in-hand with human potential. We invite all who share this ethos to join us in writing that story, together.

Applied AI

AppliedAI exists to guide humanity into the age of AI, ensuring work remains a source of meaning and boundless potential. Our mission is clear: to forge human-first AI systems that amplify our innate capabilities, addressing fundamental challenges and ensuring technology serves human advancement. We are building the essential foundation for unlocking greater human capability in the age of intelligence, liberating potential, and securing a future of shared prosperity.

opus.com