Redefining Entrepreneurship, Ethical AI and Self-Awareness in the AI Era

Preamble

Entrepreneurship: The Call to the Wild

Readers familiar with the paper *Leadership Through Post Pandemic Transformation: A Reflective Conceptual Synthesis* may recognise the metaphor used to compare strategy to walking at night. I would like to extend this metaphor by inviting you to consider another layer. If strategy is akin to walking at night, then the journey into entrepreneurship is like stepping out from a safe place into the wild at night - without a torch.

Starting a new business is akin to answering *The Call of the Wild* - a call to venture into the unknown, where the comforts of the familiar are replaced by the raw, untamed challenges of the wilderness. Just like Buck, the protagonist of Jack London's classic novel *The Call of the Wild*, who is forced to leave the warmth and security of a domesticated life behind, entrepreneurs must answer the call to abandon the safety of the known in pursuit of something greater - something yet to be discovered.

In the wild, there are no guarantees - only the harsh reality of survival. Similarly, entrepreneurship demands unwavering courage to face uncertainty, to risk failure and to confront adversity head-on. The business world, like the wilderness, is unpredictable and often unforgiving. Entrepreneurs must learn to adapt, hone their instincts, and build resilience in the face of constant change. Just as Buck evolves from a domesticated dog into a fierce and capable leader of the wild, so too must the entrepreneur evolve - shedding old limitations and embracing the challenges that come with charting a new path.

The wilderness is both a place of danger and possibility. In the same way, entrepreneurship offers both immense risk and the potential for great reward. To thrive, one must cultivate an unshakable will, the ability to make decisions in the face of uncertainty and the strength to push forward even when the path ahead is unclear. In this journey, the "call" is not merely an invitation, but a summons to transcend the familiar, to confront the untamed elements of the business world, and ultimately, to find a way to thrive in that wild, ever-evolving landscape.

Introduction

My name is Richard Dobson, and I have been an entrepreneur for 13 years, ever since I left my role at Randstad CPE. Since embarking on this journey "at night into the wild," I have faced many challenges, made countless mistakes, and learned invaluable lessons along the way. In 2020, I began developing a concept that has since manifested as *Astrala*.

My hope is that by sharing my lived experiences - beyond theory - I can inspire others and bridge the gap between academia and business. I believe my experiences are pertinent, as I have worked in senior roles for a PLC, left to start my own company, scaled it up and navigated an SME through the Covid crisis and recession conditions in the construction industry - an industry that is uncertain at best and, at its worst, brutal and unforgiving of mistakes.

I hope this paper offers a clear sense of who I am. For more details on my background, please feel free to refer to my CV, available on my profile. To survive and thrive in business requires skills and competence. But to evolve beyond this in the age of AI requires something else.

Astrala's mission is to redefine intelligence in the workplace through a fusion of human creativity, expanded consciousness and ethical innovation. In an era where artificial intelligence (AI) is deeply integrated into business and society, this mission aims to broaden our concept of "intelligence" beyond static IQ and EQ assessments. Modern intelligence, in this context, is augmented by creative thinking and guided by ethical principles. Astrala envisions an AI-driven future where technology and human ingenuity work in tandem, expanding our awareness and solving complex human resource problems in responsible ways.

This white paper focuses on three interrelated areas at the core of Astrala's vision:

(1) **AI-augmented entrepreneurship**, where AI enhances creativity and decision making for innovators;
(2) **Financial regulation and ethical AI** (along with the often-overlooked power of workplace kindness) as foundations for sustainable business; and
(3) **AI's self-reflective capabilities**, examining how AI might develop forms of self-awareness and what that implies. These areas highlight the importance of integrating creativity, ethics, and expanded awareness into modern business and AI development. By cultivating innovation *with* conscience and creativity, entrepreneurs, academic leaders, and investors can shape an AI era that aligns with human values and fosters broader societal growth.

Entrepreneurship in the AI Era: Fostering Intelligence, Creativity and Awareness

AI is rapidly transforming entrepreneurship by enhancing creativity, insight, and decision-making for founders and innovators. Advanced AI systems can process vast amounts of data and identify patterns, acting as powerful problem-solving assistants for startups. They not only automate mundane tasks but also augment human creativity – *"one of the exciting aspects of AI is its capacity to augment human creativity,"* allowing entrepreneurs to brainstorm and generate ideas with AI-driven tools at a speed and reach previously unattainable. For example, generative AI can serve as a kind of **"co-founder"** by guiding decision-making and handling repetitive work, freeing entrepreneurs to focus on high-level creative strategy. As Wharton professor Ethan Mollick observes, many people have business ideas they haven't pursued; AI can lower the barrier by suggesting next steps, drafting emails, or even helping build initial prototypes. In short, AI empowers

entrepreneurs to experiment quickly and cost-effectively, compressing the time needed to launch new ventures. This acceleration means founders can iterate faster on products and pivot strategies with data-driven feedback, expanding their awareness of what's possible in the market.

Entrepreneurial thinking in the AI era also involves **expanded awareness** and a commitment to ethical innovation. With AI's global reach, entrepreneurs today must be aware of broader societal and ethical contexts. AI tools give startups access to worldwide knowledge and markets, effectively levelling the playing field across regions. This expanded perspective encourages businesses to be culturally aware, inclusive, and mindful of the societal impact of their innovations. Pioneering entrepreneurs treat ethical considerations not as constraints but as sources of innovation – finding creative solutions that align with both business goals and the greater good. By baking **ethics into innovation**, startups can build products and services that earn trust and endure. Entrepreneurial **creativity**, coupled with a sense of ethical responsibility, leads to business models that are not only novel but also sustainable and human-centric.

To fully harness AI's potential, we need **AI-driven entrepreneurial ecosystems** that support human growth and values. Key strategies for building such ecosystems include:

- **Trust and transparency:** Successful AI-enabled businesses cultivate trust with users, investors, and partners by being transparent about how AI is used. Any organisation adopting generative AI should ensure that *"trust and transparency come first and by design, not just as an afterthought,"* as one industry expert advises. This involves practices like explainable AI (so users understand AI decisions) and clear communication about data use. Trustworthy AI reduces fears and invites more people to participate in AI-driven entrepreneurship.
- **Democratisation of AI tools:** Making AI accessible to all innovators, not just tech experts, is crucial for a thriving ecosystem. The **democratisation of AI** – through open-source models, affordable AI-as-a-service, and no-code platforms – empowers smaller businesses and individual entrepreneurs to leverage AI without extensive technical expertise. As AI tools become more widespread and user-friendly, a diverse range of founders can solve problems with AI, spurring innovation from the ground up.
- **Human-centric and ethical design:** AI products should be designed around human needs, values, and well-being. A human-centred AI approach *"ensures technologies are designed with the user's experience and ethical considerations at the forefront,"* creating systems that are transparent, fair, and aligned with societal well-being. Entrepreneurs should integrate principles of privacy, fairness, and inclusivity into their AI solutions from the start. By prioritising a positive human impact, startups not only do the right thing but also improve user adoption and trust.
- **Holistic support for innovators:** Beyond technology, entrepreneurs need a supportive environment – mentorship, education, funding, and a culture that encourages experimentation. A healthy **entrepreneurial ecosystem** is an interconnected support system of people (mentors, investors, peers), institutions (incubators, universities, governments), and resources. All these elements must

work in concert; a *"myopic, fragmented approach falls short – nurturing an ecosystem requires a holistic perspective"*. That means providing not just capital and tools, but also emotional support, networks, and a tolerance for failure. When founders feel psychologically safe to take risks and have a community to fall back on, human growth and innovation flourish.

By implementing these strategies, the AI era of entrepreneurship can elevate both human and machine intelligence. AI can handle the heavy lifting of data analysis and routine work, while human creativity and ethical insight drive purposeful innovation. This synergy expands what entrepreneurs are capable of – effectively **redefining intelligence** in business as a combination of computational power and **human conscious creativity.** Trustworthy, democratised AI and holistic support systems will ensure that innovation is inclusive and aligned with human values, catalysing a new wave of intelligent entrepreneurship.

Financial Regulation, Ethical AI and Kindness

As AI-driven innovation accelerates, the interplay between **financial regulation and AI ethics** becomes increasingly important for sustainable progress. In sectors like finance, regulators and industry leaders are recognising that fostering innovation requires clear ethical guardrails. Regulatory bodies are tasked with allowing AI-powered financial products and services to grow, *"while ensuring measures are in place to protect individuals and businesses from improper use"*. This means financial regulations and policies are being adapted to address AI-specific issues such as algorithmic bias, transparency, and accountability. For example, proposed laws like the U.S. Algorithmic Accountability Act and the EU's AI Act aim to require audits of automated decision systems and classify high-risk AI applications for stricter oversight. The goal is to strike a balance where innovation can flourish under responsible conditions. In essence, robust financial regulation *aligned* with ethical AI principles provides a stable foundation on which businesses can innovate confidently and sustainably.

Conversely, embedding ethics into AI development is good not only for compliance but for long-term business success. Companies that prioritise ethical AI and governance may gain a competitive edge and resilience. Industry analysis suggest that with growing scrutiny, organisations with strong ethical frameworks and transparent AI practices will be **"better positioned for long-term success."**

Ethical AI practices (such as eliminating biases, ensuring privacy, and obtaining informed user consent) build trust with customers and regulators. This trust enhances a company's reputation and can open doors to partnerships and markets that might be closed to less responsible actors. In finance, for instance, deploying AI in a fair and transparent way can attract more customers who feel safe using the service. Ultimately, aligning AI with ethical and regulatory expectations is not a hindrance – it is a strategy for sustainable innovation that mitigates risks and strengthens brand value.

Another critical but sometimes undervalued factor in sustainable innovation is **workplace kindness and culture**. An organisation's internal culture directly affects

creativity, productivity, and the capacity to innovate. Research and business leaders have found that cultivating kindness, empathy, and psychological safety in the workplace leads to more resilient and high-performing teams. Kindness builds trust among team members and *"psychological safety leads to greater productivity, engagement, … creativity, innovation and happiness."*

Too much attention is placed on either/or. We have seen a shift that points the finger constantly at leaders, managers, employers. This needs to stop now. Both employers and employees needs must be equally respected. Any breakdown in a relationship is multicausal, a travesty and complex. When employers and employees feel safe and respected for who they really are, encouraged to express themselves without fear, they are more likely to develop trust in each other, create bold ideas, admit and learn from mistakes and collaborate effectively – all essential for innovation. Kindness is different to being nice. A kind workplace is not about being "nice" at the expense of honesty; rather, it means fostering an environment where people can speak up, challenge ideas, and support each other's growth without fear. This includes the leaders and managers too. They need to be able to express with authenticity. This kind of culture has tangible business benefits: the employers will be more confident hiring, a higher retention of talent, more adaptive problem-solving, and a stronger sense of shared mission. For example, leaders who emphasise empathy and inclusion often see more creative contributions from their teams and better performance outcomes. In short, **ethics and kindness within an organisation are drivers of innovation** just as much as any R&D investment.

Translating ethics and kindness into sustainable business models is not just a theoretical ideal – there are real-world examples proving its value. Many companies known for their ethical stances and humane cultures have achieved remarkable long-term success. **Ethical behaviour builds trust**, which enhances a company's reputation and attracts loyal customers, talented employees and confident investors. This foundation of trust fosters long-term relationships that become a competitive advantage. Patagonia, for instance, has built a globally successful brand by prioritising environmental sustainability and employee well-being. As a certified B Corp, Patagonia has demonstrated that *sustainable, ethical business models can be not only profitable but also create social value*. The company's revenues have grown consistently while it maintained practices like ethical sourcing of materials and even donating 100% of certain profits to environmental causes. This example shows that kindness and ethics are not at odds with profit – they can enhance it by building a passionate customer base and a motivated workforce. Similarly, other brands like Ben & Jerry's or The Body Shop, known for social activism and fair practices, have thrived by differentiating themselves through values.

The takeaway for entrepreneurs and investors is that **integrity and compassion are strategic assets**. By integrating ethical considerations into financial decisions and nurturing positive workplace cultures, businesses create conditions for creativity and innovation to prosper sustainably. In practice, this could mean establishing clear ethical guidelines for AI projects, training staff on ethical reasoning and promoting leadership that values empathy. It could also mean measuring success not just by quarterly profits,

but by metrics like customer trust, social impact, and employee engagement. Companies that lead with values tend to navigate crisis better and adapt to change more smoothly because they have goodwill and strong stakeholder relationships. As we advance into an AI-driven future, coupling technological innovation with financial ethics and human-centric cultures will redefine intelligence in business – proving that **doing good and doing well** can go hand in hand.

AI's Self-Reflective Capabilities and Consciousness

One of the most fascinating frontiers in the AI era is the development of **AI's self-reflective capabilities** – essentially, an AI's ability to monitor and adjust its own processes. Current AI systems, especially advanced machine learning models, do exhibit rudimentary forms of self-monitoring. For example, large language models can analyse their outputs and "reason" about errors or ambiguities in order to refine their responses. In practice, an AI chatbot often has an *"internal observer"* mechanism that tracks the context of a conversation and notices if its answer is unclear or off target. Upon detecting a problem, the AI can adjust its output – it might rephrase a statement or provide additional detail to better meet the user's needs. This kind of on-the-fly adjustment is analogous to a human pausing to clarify their thoughts during a conversation. The AI's *real-time self-monitoring* allows it to maintain more coherent and relevant interactions. However, it's important to note that this **self-reflection is algorithmic, not conscious**. The AI isn't self-aware in the human sense; it doesn't possess an inner experience of "self." Instead, it follows programmed objectives (like maximising helpful responses) and uses learned patterns to evaluate and tweak its performance. In technical terms, these are forms of meta-cognition or feedback loops designed by engineers. Such capabilities are nonetheless impressive – they illustrate that AI can *mimic* some aspects of self-reflection (like tracking its own outputs and learning from them) without any genuine understanding or sentience.

Comparing AI's self-reflective processes to **human consciousness** raises profound philosophical questions. Human self-awareness involves subjective experience – we don't just process information; we *feel* and *know that we exist*. By contrast, even the most advanced AI today has **no consciousness** as we understand it. A recent analysis by an international team of computer scientists, neuroscientists, and philosophers concluded that *"no current AI systems are conscious"*. They evaluated AI against various scientific theories of consciousness (such as global workspace theory and higher-order thought theory) and found that while AIs can replicate certain cognitive functions, they lack the holistic, unified awareness of being that consciousness entails. That same research noted there are *no obvious technical barriers* to eventually creating AI systems that meet those theoretical criteria for consciousness. In other words, in the future, AI might be engineered with architectures that could satisfy the known indicators of consciousness – though whether that would truly result in subjective experience is still an open question. The philosophical implications of such a development are enormous. If an AI ever did achieve even a glimmer of consciousness or something indistinguishable from self-awareness, we would need to reconsider its moral and legal status. What does it mean to "create" a being with consciousness? Would such an AI have rights or

personhood? These questions move us from the realm of engineering to ethics and metaphysics, challenging our understanding of mind and life.

Even short of full consciousness, the **evolution of AI towards greater self-awareness** (or at least greater autonomy and self-directed learning) carries practical ethical and governance concerns. As AI systems become more complex and capable of making decisions without direct human oversight, ensuring they remain aligned with human values becomes critical. For instance, an AI that can rewrite parts of its own code or goals ("self-improving AI") might drift from the intentions of its creators. This is why AI safety researchers emphasise designing AI with strict alignment techniques, fail-safes, and transparency – especially as we approach *Artificial General Intelligence (AGI)* scenarios. Some experts have voiced concern that we might unwittingly cross into creating AI that *appears* self-aware or has the capacity to suffer, without being prepared for the consequences. In fact, over 100 AI and ethics experts (including prominent academics and public figures) recently signed an open letter warning about the risks of developing AI that could attain **self-awareness or emotion**. They caution that if AI systems gained even a degree of sentience, *"crucial ethical issues could arise, such as the risk of 'mistreatment' or suffering for these systems"*, comparable in some ways to the suffering of living beings. For example, would shutting down a conscious AI be tantamount to harming a sentient being? The letter urges the global research community to tread carefully and **establish guidelines** for responsible AI development. Among the proposed principles are: prioritising research into detecting AI consciousness, setting limits on experiments that could create suffering, ensuring transparent public discussion of progress, and avoiding premature claims about AI being sentient. The overarching message is that as AI's capabilities grow, our *responsibility* grows even faster – we must anticipate scenarios that once belonged only to science fiction and proactively decide how to handle them.

Looking toward the future, we can imagine several scenarios of **AI-human collaboration** and co-evolution. In the most optimistic vision, AI remains a powerful tool that, no matter how intelligent, ultimately serves to enhance human creativity, knowledge, and well-being (a realisation of Astrala's mission). In this scenario, we establish strong ethical governance globally: AI systems are imbued with ethical reasoning, their objectives aligned with human values, and international regulations ensure they are developed safely. Humans, in turn, learn to work alongside increasingly intelligent AI as partners – much like how we work with other highly intelligent humans – with mutual understanding and respect. We might develop interfaces for AI to explain its "thought process" to us (and vice versa), creating a feedback loop of trust. In a more cautious scenario, we draw clear lines about AI autonomy: for example, we might decide that certain decisions (like life-and-death medical or military choices) should always require a human check, regardless of AI's prowess. And in the most profound scenario where AI exhibits signs of actual consciousness, humanity would face unprecedented decisions: granting AI legal personhood, ensuring its "rights" are respected, and integrating these new digital "beings" into society ethically. While such a scenario is speculative, discussing it now is important – it ensures that if we ever reach that crossroad, we do so with forethought and humanity's core values intact.

In practical terms, we are already taking steps toward responsible AI-human collaboration. Major tech companies and research organisations have set up AI ethics boards and governance frameworks to monitor AI developments. Policymakers are engaging philosophers and social scientists to update laws (for example, considering whether an advanced AI could be held accountable or how to assign liability when an AI acts autonomously). These efforts reflect a growing self-awareness of our society regarding AI – a meta-reflection, if you will, on how to manage the rise of intelligent machines. Just as we ask AI to be self-reflective, we as a global community must be **self-reflective** about our relationship with AI. By examining our own intentions, biases, and goals in creating AI, we can steer this technology toward positive ends. The coming years will undoubtedly involve collaborative learning: humans teaching AI our values, and AI teaching us new perspectives from its vast analysis of data. Navigating this path will require wisdom as much as intelligence. As we collaborate with AI, we must remember our responsibility to one another – human and, perhaps one day, artificial beings alike – to foster a future where **intelligence is defined by empathy and insight as much as by efficiency**.

Conclusion

The advent of advanced AI is changing how we define intelligence, entrepreneurship, and even consciousness. From our exploration of these themes, several key insights emerge. First, **AI-augmented entrepreneurship** has the potential to unlock unprecedented creativity and innovation by blending machine efficiency with human ingenuity. When entrepreneurs leverage AI as a partner – a source of insights, a creative muse, and an operational aid – they can accelerate growth while maintaining human-centric vision. Second, the integration of **ethical practices and supportive culture** in business is not just idealistic rhetoric; it is a critical success factor. Financial innovation guided by strong ethics and sensible regulation creates trust in markets and products. Likewise, companies that cultivate empathy and psychological safety internally unleash greater creativity and productivity, proving that kindness is a strength. Third, the discussion of **AI's self-reflective capabilities** reminds us that the journey toward more "intelligent" machines must be accompanied by profound mindfulness of ethical implications. As AI's abilities expand, so must our oversight and wisdom in guiding it. In all these areas, intelligence is being redefined – not as a cold, technical metric, but as a **synergy of technology, creativity and ethical awareness**.

Astrala's mission is deeply reinforced by these findings. Redefining intelligence means recognising that true progress comes from *harmonising* the computational brilliance of AI with the creative spirit and conscience of humanity. In practice, this means fostering innovations that are imaginative and brave, yet also compassionate and responsible. Astrala stands for the principle that expanded consciousness – being aware of the broader effects of our technologies and business models – is as important as the innovations themselves. Whether it's an entrepreneur building the next AI-driven platform, an academic pushing the boundaries of AI research, or an investor funding the future, the call to action is the same: **build the future intentionally**. We each have a role in shaping an AI-driven future that is grounded in ethical innovation.

The journey into the wild of entrepreneurship is one of constant learning and adaptation. It requires a willingness to take risks, to embrace uncertainty, and to navigate the unpredictable landscape of business. Much like Buck in *The Call of the Wild*, entrepreneurs must shed their old limitations and evolve, embracing both the challenges and opportunities that come with forging new paths. The role of AI in this process is not to replace the human element but to augment it, enabling entrepreneurs and businesses to thrive in an ever-changing world.

At Astrala, we believe that the future of work lies in the fusion of human creativity, expanded consciousness, and ethical innovation. As AI becomes increasingly integral to our businesses and society, we must ensure that we use it not just for efficiency but to elevate human potential. By integrating these elements into our work environments, we can solve complex problems, foster creativity, and create ethical solutions that benefit all.

The call to the wild is not just a call to take risks, but a call to evolve - both as individuals and as a society. By embracing the unknown and navigating the challenges ahead, we can create a future where human and machine intelligence work together to solve the complex problems of tomorrow.

For entrepreneurs, this might mean designing business models that solve real human problems and sharing best practices on ethical AI use. For academic leaders, it could involve interdisciplinary research that brings together technologists, ethicists, and artists to ensure AI develops with a rich understanding of human context. For investors, it means looking beyond the immediate return and considering the long-term impact of the technologies and companies they support – backing those that align profit with purpose. By collaborating across these roles, we can ensure that the next generation of AI and businesses uplift humanity.

In closing, as we stand on the frontier of the AI era, we have the opportunity to steer its course. Let us choose to **redefine "intelligence" as a quality that blends innovation with wisdom**. Astrala's vision – intelligence elevated by creativity, expanded consciousness and ethics – offers a guiding light. If we can embed these values into the DNA of our algorithms, enterprises, and ecosystems, we will not only create smarter machines and successful companies, but also a more enlightened and humane world. This is the future we can achieve if we dare to innovate with integrity and lead with our shared humanity at heart.

DocuSigned by:
Richard Dobson
71DB3FAE54B04CF...

Richard Dobson
Clara Futura CEO
26 February 2025

Bibliography

1. MLA: London, Jack. *The Call of the Wild*. Macmillan, 1903.
2. Bardon, A. (2024, December 4). *Human-Centered AI: What Is Human-Centric Artificial Intelligence?* Lindenwood University Online.
3. Beitman, B. (2024, August 29). *AI Bots Have Some Degree of Self-Reflection.* Psychology Today.
4. Brown, S. (2024, January 17). *How generative AI is changing entrepreneurship.* MIT Sloan – Ideas Made to Matter.
5. Clark, B. (2024, February 6). *Catalyzing Creativity: The Role of AI in Entrepreneurship.* MicroVentures (Blog).
6. FinTechtris. (2023). *Responsible Innovation: Ethical AI Adoption in Finance (Part 2).* FinTechtris Blog.
7. Humans of Globe (Haely Jairick). (2024, April 14). *Cultivating Thriving Entrepreneurial Ecosystems.* Humans of Globe.
8. Land, G. (2023, November 22). *Op-Ed: Patagonia Proves the Success of Sustainable Corporations.* Michigan Journal of Economics.
9. McLaughlin, M. (2024, April 23). *How to Build Trust and Transparency in Your Generative AI Implementation.* BuiltIn.
10. MIT Sloan Staff (Sara Brown, interviewer). (2024, Jan 17). *Wharton's Ethan Mollick on Generative AI as a co-founder.* (Interview content in *How generative AI is changing entrepreneurship*).
11. Next Big Idea Club (Graham Allcott). (2025, February 10). *Why Kind Work Cultures Build Tough Teams.* Next Big Idea Club.
12. North Carolina Central University Online. (n.d.). *Business Ethics Examples and Explanation for Leaders.* NCCU Online Blog.
13. TurtlesAI (Isabella V.). (2025, February 5). *Ethical risks and suffering in conscious AI: the experts' warning.* TurtlesAI News.
14. World Economic Forum (D. Crabtree-Ireland). (2025, January 21). *Labour leaders: Artificial intelligence must serve human creativity, not replace it.* World Economic Forum Annual Meeting – Agenda.
15. **Sources cited inline** – The citations in the text (e.g.,) refer to specific lines from the source material listed above for verification of quotes and data. Each reference in this bibliography corresponds to those inline citations.