
Thought Piece:

Always Read The Manual

Tradecraft

What do the unlikely trio of Messrs Kirby, Kamprad and Hastings all have in common, apart from being exceptionally successful business founders and managers? They all chose to divulge the secrets of their companies' success in the form of a 'manual'. Tom Kirby's Little Red Book, Inka Kamprad's The Testament of a Furniture Dealer and Reed Hastings's Culture Deck all laid-out, with exceptional clarity, the building blocks of their businesses' enduring and world-class culture. Each did it early in the lifetime of their business', too. Notwithstanding that all three were primarily motivated to inspire employees rather than competitors, the apparent revelation of the secret sauce of these businesses does not seem to have eroded their competitive moats one iota. An observer might reasonably ask: why were these businesses not cloned?

Eating the world

In 2011, not long after the last of the aforementioned manuals were shared, Marc Andreessen published his now-prophetic Wall Street Journal op-ed 'Software Is Eating the World'. Looking back, Andreessen describes a world which today almost seems quaint. The article also put down a marker on what was, in hindsight, the beginning of a new era of accelerated codification of business and work practices. Today he might title the article, 'The Agent is Eating the World'. It is fair to assume that those manuals are by now well ingested.

The plethora of AI agents emerging today still need spoon-feeding from us humans. Everyone knows LLMs need context, but one cannot overstate that point. Models are extremely powerful at predicting 'what next' at the token level, but they are also, in effect, uber-generalists, so they need granular direction and feedback to pick up on the nuances of a complex real-world task. AI Models are extremely good at linear/repetitive algorithms, and their abilities in lateral tasks are improving.

One might offer three basic layers of abstraction to describe and help predict the real world: mental models (the ontology), codification (the taxonomy/rules) and algorithms (the implementation). Henry Ford and Aldred Sloan began codifying work in the mid-20th century as they sought consistent and repeatable factory workflows to drive efficiencies at scale. The first AI agents, chatbots, have a similar linear approach. What is new and exciting today is the promise that more advanced agents hold, to eventually codify and automate the more abstract layers of businesses, especially human decision making.

For an agent to replicate high-level human decision making, requires breaking down the tacit or embedded knowledge that underpins those decisions. Such knowledge can be hard to identify, never mind automate. Tacit knowledge refers to knowledge that is understood or employed without it being fully expressed or documented in spoken or written words. It can be knowledge that sometimes is just 'felt' and when you try to put it into words, the words never fully capture the meaning. It might also be thought of in the context of acquired mastery, good taste or what is called discernment. Anyone who has tried to put words to the subtleties of a great wine, a piece of art or even an advanced yoga pose will also relate.

No bad language please

So, in the expression and understanding of real-world actions and decisions for feeding into LLMs, the words employed matter a lot. How can you and I debate something if we have different interpretation of the words being used? Businesses seeking to standardise their work methods already understands this; Toyota employs a common engineering language to

ensure that all definitions and specifications mean the same things to all involved throughout its supply chain. At the other extreme, Palantir today talks of a "full-scale, full-fidelity semantic representation of the enterprise". Somewhere in-between, here at Phoenix, a colleague building an advanced agent for our 'DREAM' investment framework recently noticed his agent's at-times inconsistent interpretation of the common business phrase "pricing power" (since cleverly resolved – nice work, Dan!).

Palantir has been at the forefront of efforts to codify and automate workflows in governments, armies and old school industries. The company pioneered the concept of 'the Forward Deployed Engineer' – in effect, a paid-for Trojan horse that deploys business-minded 'solutions engineer's to intimately and fastidiously learn about the nuances of a client's business practices (and thus the market vertical that it operates in). In doing so, the FDE systematically codifies that company's embedded knowledge through inquiry and bespoke systematic learnings loops. Palantir then banks and compounds that knowledge via its proprietary operating system (for resale!). Infamous for many of its partnerships with government and defence agencies, Palantir has genuinely pioneered some groundbreaking approaches to codifying work for AI, not least through what it refers to as its Ontology. Here is how it describes it:

The Ontology represents the decisions in an enterprise, not simply the data... (it) integrates this data into a full-scale, full-fidelity semantic representation of the enterprise...

(it) safely captures the decision data produced by operational users as they carry out daily work...

(and it) natively models actions within a cohesive, decision-centric model of the enterprise.

(emphasis ours)

Palantir's Ontology is also referred-to as an organisation's 'digital twin' – a live, dynamic, two-way model of the enterprise. It sounds like a digital manual!

The last liberal art

Which brings us to investing, where another master, Mr. Buffett, also goes to great lengths to openly share his tradecraft in a manual of sorts: his annual investor letters. If LLMs are on a path to codify the tacit knowledge required to predict complex work patterns and human decisions, then surely fundamental investing could be similarly ripe for disruption.

Howard Marks once wrote that there are two things that can be the source of superior investing: a better comprehension of the future and a superior ability to process qualitative information. As examples of the qualitative information, Marks offered 'market mood', 'quality of management', 'effectiveness of a company's product development' and the 'strength of its accounting'. All are, to varying degrees, the types of abstract knowledge that requires a form of mastery and discernment to assess properly. (And all are areas of our framework that we are actively augmenting with AI at Phoenix.)

Most of Marks's qualitative categories, involve assessing human behaviour, which is very hard, given that humans are nothing if not inconsistent. Charlie Munger, the great lateral thinker, recognised this and spent a lot of his life learning deeply about human psychology. Munger also popularised the term 'mental models'; what he was really trying to do, in effect, was to

codify at the highest level of abstraction the elementary basics of successful business and management.

Munger also emphasised the importance of integrity in life and business. His and Buffett's networking off-sites (the so-called 'Graham Group') were a key component of their creative debate and due diligence on ideas and people (if only we had the minutes of those meetings!). In the quest to assess management integrity, in-person assessment, trusted references and testimonials – i.e. human oriented assessment – has long been a preferred approach. Yes, one can invert and perhaps codify the search for obvious signs of low-integrity but the problem is that human behaviour is so complex to assess and capable of masking one's true self. Some of us are acutely wired to sense 'flakiness' in others and gut feelings are probably relied-upon more than is admitted. The fact that elites like Buffett, Dimon and Ackman were each, on occasion, very publicly on the wrong side of a judgement of a colleague's integrity shows how difficult it is to assess. AI agents can augment an assessment by widening the net of source material and helping identify patterns but for now it seems that assessing humans' integrity will remain an augmented task that warrants a flexible approach.

Mental flexibility

In May 1969, when Buffett was 39 years-old and announcing the closure of the Buffett Partnership, his letter to his investors raised the question whether he had enough mental flexibility to cope with the "frustrating and negative" investing environment at that time. He also wryly relayed an observer's comment that perhaps security analysts over 40 "know too many things that are no longer true"¹.

In that spirit, we humbly suggest that such mental flexibility is vital today in order to be open to the idea that AI agents are eating into the world of business and investing but also, to retain the immense value of human discernment.

(¹ and to answer the obvious question: by comparison, the average age of the Phoenix investment team is only 39.6!)

This document is issued by Phoenix Asset Management Partners Limited (PAMP), registered office 80-82 Glenthams Road London SW13 9JJ, Company number 03514660. Authorised and regulated in the UK by the Financial Conduct Authority. This advertisement is directed only at persons in the UK who are (i) investment professionals as defined in Article 19 of the Financial Services and Markets Act 2000 (Financial Promotion) Order 2005 or (ii) are persons of a kind described in Article 49(2) of the Promotion Order or (iii) are persons to whom this presentation may otherwise lawfully be issued or passed on.

Castelnau Group Limited ("CGL") is a Guernsey closed-ended company listed on the Specialist Fund Segment of the London Stock Exchange. Shares traded on a stock market will fluctuate in accordance with supply and demand and may not reflect the underlying net asset value of the shares. The value of investments and any income from them may go down as well as up and investors may not get back the amount invested. There can be no assurance that CGL's investment objective will be achieved, and investment results may vary substantially over time. Past performance is not a reliable indicator of future performance. This advertisement is for information purposes only and does not constitute an offer or invitation to purchase shares in CGL.

An investment in CGL involves a high degree of risk, including the risk that the entire amount invested may be lost. Buying a concentrated portfolio of both listed and unlisted companies can result in above average volatility. Investors should therefore be open to some risk and volatility. There can be a long lag between us making an investment and its value being reflected, so CGL is intended for investors with a long-term investment horizon. Although investors can sell shares on any day, if your horizon is short-term then our approach may not be a suitable home for your money. CGL is intended for professional clients and “sophisticated” or “High Net Worth” retail investors, those considered as “advanced” but is not suitable for retail investors with only basic knowledge of investments.

The information contained within this document has been obtained from sources believed to be reliable and accurate at the time of issue. Prospective investors should not rely upon this document for tax, accounting or legal advice and should consult their own advisors prior to making any investment. The Prospectus and other regulatory documents can be found at www.castelnaugroup.com

Contact Details

Castelnau Group Limited

www.castelnaugroup.com

Investment Manager

Phoenix Asset Management Partners Ltd
80-82 Glenthams Road
London
SW13 9JJ
Tel: +44 (0)20 8600 0100
Email: Info@castelnaugroup.com

Phoenix is authorised and regulated by the Financial Conduct Authority
Registered as an Investment Adviser with the Securities and Exchange Commission

Administrator

Northern Trust International Fund Administration Services (Guernsey)
Les Banques, Trafalgar Court
St Peter Port
GY1 3QL