LLOYD TEVIS
INVESTMENTS, LLC

Securing your future through Precision Investing™

# Research Note Series

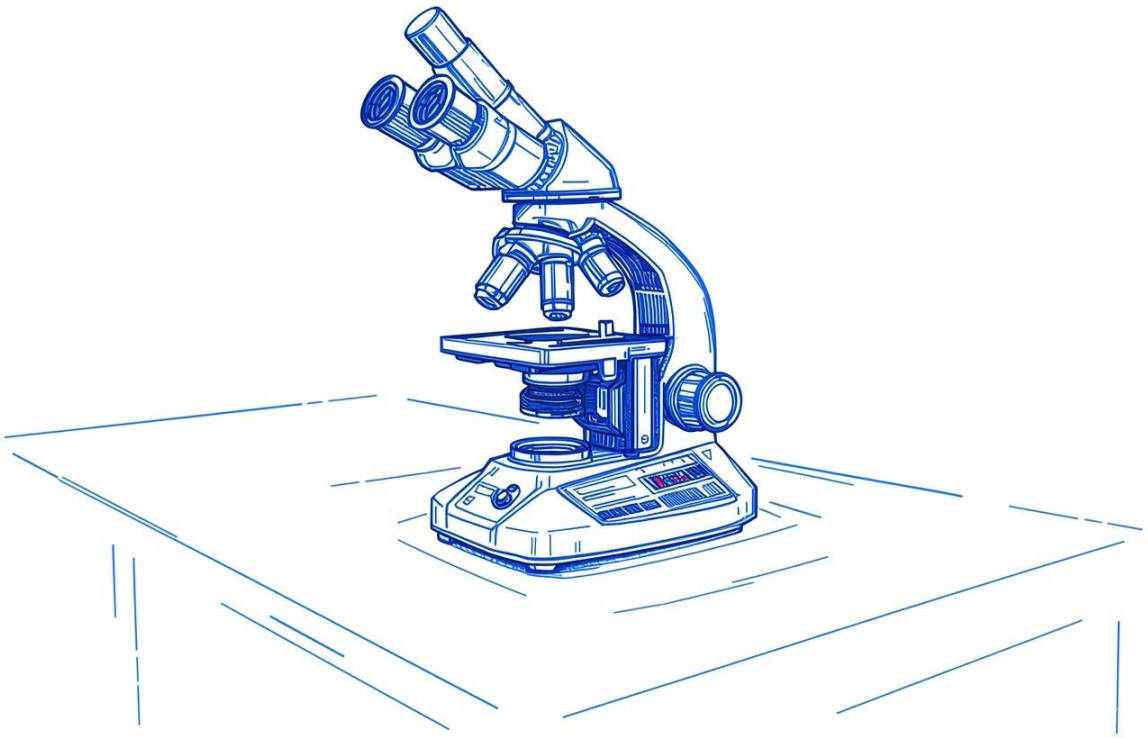

# Recognizing Systemic Risk In Applications Of AI

April 2026

It is hardly news to say AI has sparked great interest. But we have been here before – we need not go blindly forward. About 150 years ago power tools began to be introduced into manual labor occupations. The near limitless application of mechanical advantage was transformative. What could be done was vastly expanded, great business opportunities were opened up, much back breaking work no longer needed to absorb tens of thousands of lives and cause death or injury to a percentage of them, and incomes of the tool operators soared well above that of the former laborers. The hope is that AI will bring similar positive changes to office work – long a sink of many millions of hours of human time and diligence.

But how to do it? Let us look back to the earlier transformation for insight. A naive idea on how to apply mechanical advantage to the dreary task of ditch digging might look like this

But in fact we know that the right answer to this problem initially looked like this

And after 120 years of refinement it now looks like this.

There are two points to take away. First the naive idea is wrong. Second not that much changed between the original steam shovel and the modern power shovel. In other words, the first non-naive idea was pretty close to the optimal idea.

When we compare the steam shovel to the exoskeleton it is easy to identify its sources of advantage:

1. Greater operator safety and comfort
2. Fewer moving parts leading to greater reliability and lower maintenance cost
3. Much better ability to scale the solution to the size of the task

In short, system engineering analysis shows why the steam shovel was the way forward and not the exoskeleton.

When we look at AI we see that many current applications are closer to exoskeletons than to steam shovels. There are I think three reasons this is so. First there is some existing process in place. An engineer tasked with injecting mechanical advantage will first look at the points in the process where his technology can be applied.

Second, there are people doing the process and it is natural to first think of ways to make the people more powerful. These are both naive thoughts and they lead to exoskeleton solutions. In other words to an incorporation of the technology which keeps the current worker in place but simply tries to help him to do more – to be more productive.

There is third reason that is perhaps AI specific. The technology is not entirely reliable yet. The technology hallucinates or just goes down the wrong path a percentage of the time. Those of us who work with the raw technology have observed impressive improvement in the past year. We have a confidence that the technology will soon be as reliable as people. But until it is, we see a need to keep humans in the loop checking its work, offering guidance and capable of taking over as a safety driver. These considerations suggest keeping the application of AI in the exoskeleton mold rather than the steam shovel mold.

But the challenges of this approach are showing up. Whereas before a desk worker sat down and did his job, now he is half doing his job and half supervising a team of enthusiastic but not totally reliable AI collaborators. Worse yet some of the collaborators may go through abrupt personality shifts as their underlying models are upgraded. Operator fatigue is showing up. And the idea of a safety driver is proving a cruel myth. Once one builds the assumption of mechanical advantage into job tasking, client commitments and financial plans one cannot fall back to the manual way of doing things without massive disruption.

In short we are seeing the system issues of operator safety and comfort, reliability, maintenance and efficient task sizing showing up.

How do we avoid falling into the exoskeleton trap? Instead of focusing on the worker or the process, I think we have to focus on the task to be accomplished. If we ask how do "I empower ditch diggers" we will never imagine steam shovels. If we ask "How do I pick up dirt from here and plop it there" the steam shovel solution is almost self evident. And I think this will prove a reliable guide to deciding if a proposed solution is of the exoskeleton class or the steam shovel class. If it is process or worker centered it is an exoskeleton. If it is objective centered then it is a steam shovel.

This insight also explains why not much changed in going from the steam shovel to the power shovel. Because the basic objective did not change the concept of a solution did not change. What in fact were the changes? There were three. First the steam engine was replaced by an diesel engine driving an electric generator. This was a component upgrade which changed nothing about the solution. Second, in the original steam shovels power was applied to task via winches that pulled on wire ropes. This technology is still used in the largest machines, but in lower power tasks

hydraulic power transmission has proved to have a better cost structure and greater nimbleness. This is a refinement of a component to the needs of a particular niche. Finally the original machines moved along custom laid rails, while modern machines crawl on treads. This is adoption of a later invented component to add value to the original concept. As we see the objective did not change, the solution concept did not change and refinement basically just applied otherwise invented technologies to enhancing the concept.

Turning back to AI, we noted one argument for the exoskeleton solution to be the need to keep the human in the loop to catch the errors of the technology. We would argue that in part is due to misapplication of the technology. If you have a partially reliable technology the natural thing to do is to focus the role of the technology. Use it only where it brings key improvements. But with AI we notice the tendency is rather the opposite. People try to use it everywhere for everything. This we think is a mistake. If a batch file can coordinate a set of subprocesses that is a far more reliable way than relying on AI orchestration. If a decision can be taken by a sequence of if-then-else encoded rules, that is how the decision should be taken. It will be more reliable, higher performing, cheaper to run and faster to get running reliably than an AI solution will be. Relying on AI only where it adds essential value will reduce the supervision problem.

Currently business mangers and venture capitalists are plowing large sums into AI development. The two cases are not identical. For business managers to quickly roll out some naive exoskeleton projects to learn about the technology is perfectly reasonable. A quickly built project is most likely a limited investment that can be rapidly replaced as technological understanding grows. Venture capitalists, however, should be cautious about building a portfolio of exoskeleton projects. This could be a systemic risk factor that sinks the portfolio. In general a VC funded technical solution needs 5 to 10 years to grow to the scale that justifies the risk taking. An exoskeleton solution that is bulldozed by a power shovel in three years will prove a most disappointing VC investment. Much more so than business managers, VCs need a sharp eye to distinguish between exoskeletons and power shovels.

It is an exciting time for technologists, businessmen and investors. There is much to learn, do and accomplish. Our closing thought is an exhortation – "Engage - Thoughtfully"

About the Author
Nicolo G. Torre PhD, CFA is a noted quantitative investment manager. He has made fundamental contributions to algorithmic trading and portfolio risk management as detailed in the "The Market Impact Model" (available online) and the Wikipedia ar-

ticle"Multiple Factor Models." His strategy work has included the initial investment case for US Treasury TIPS and the first ETF investment strategies for IShares and State Street Global Advisors. This work supported the introduction of new security types which have grown to $7 trillion in outstanding (TIPS) and $15 trillion in outstanding (ETFs.) In the domain of retail investing Torre contributed the investment process to the first RoboAdvisor (Sharebuilder.) Roboadvisors currently manage 5% of retail assets (about $2 trillion domestically.) Torre also developed the concept of goal based investing. This concept has become ubiquitous in how investment firms engage with retail clients. Most recently Torre is the Founder and CEO of Lloyd Tevis Investments. Lloyd Tevis is the first firm to offer an autonomous AI investment adviser. The adviser implements goal based investing within a RoboAdvisory business model. As an expert in both portfolio risk management and in applying AI to re-engineer investment processes, Torre brings a unique perspective to the topics discussed in the current research note.