Vol. 1, No. 2 April 2018

New Machinist Journal

Real Time Learning: A Better Approach to Trader Surveillance

Dr. Jimmie Lenz

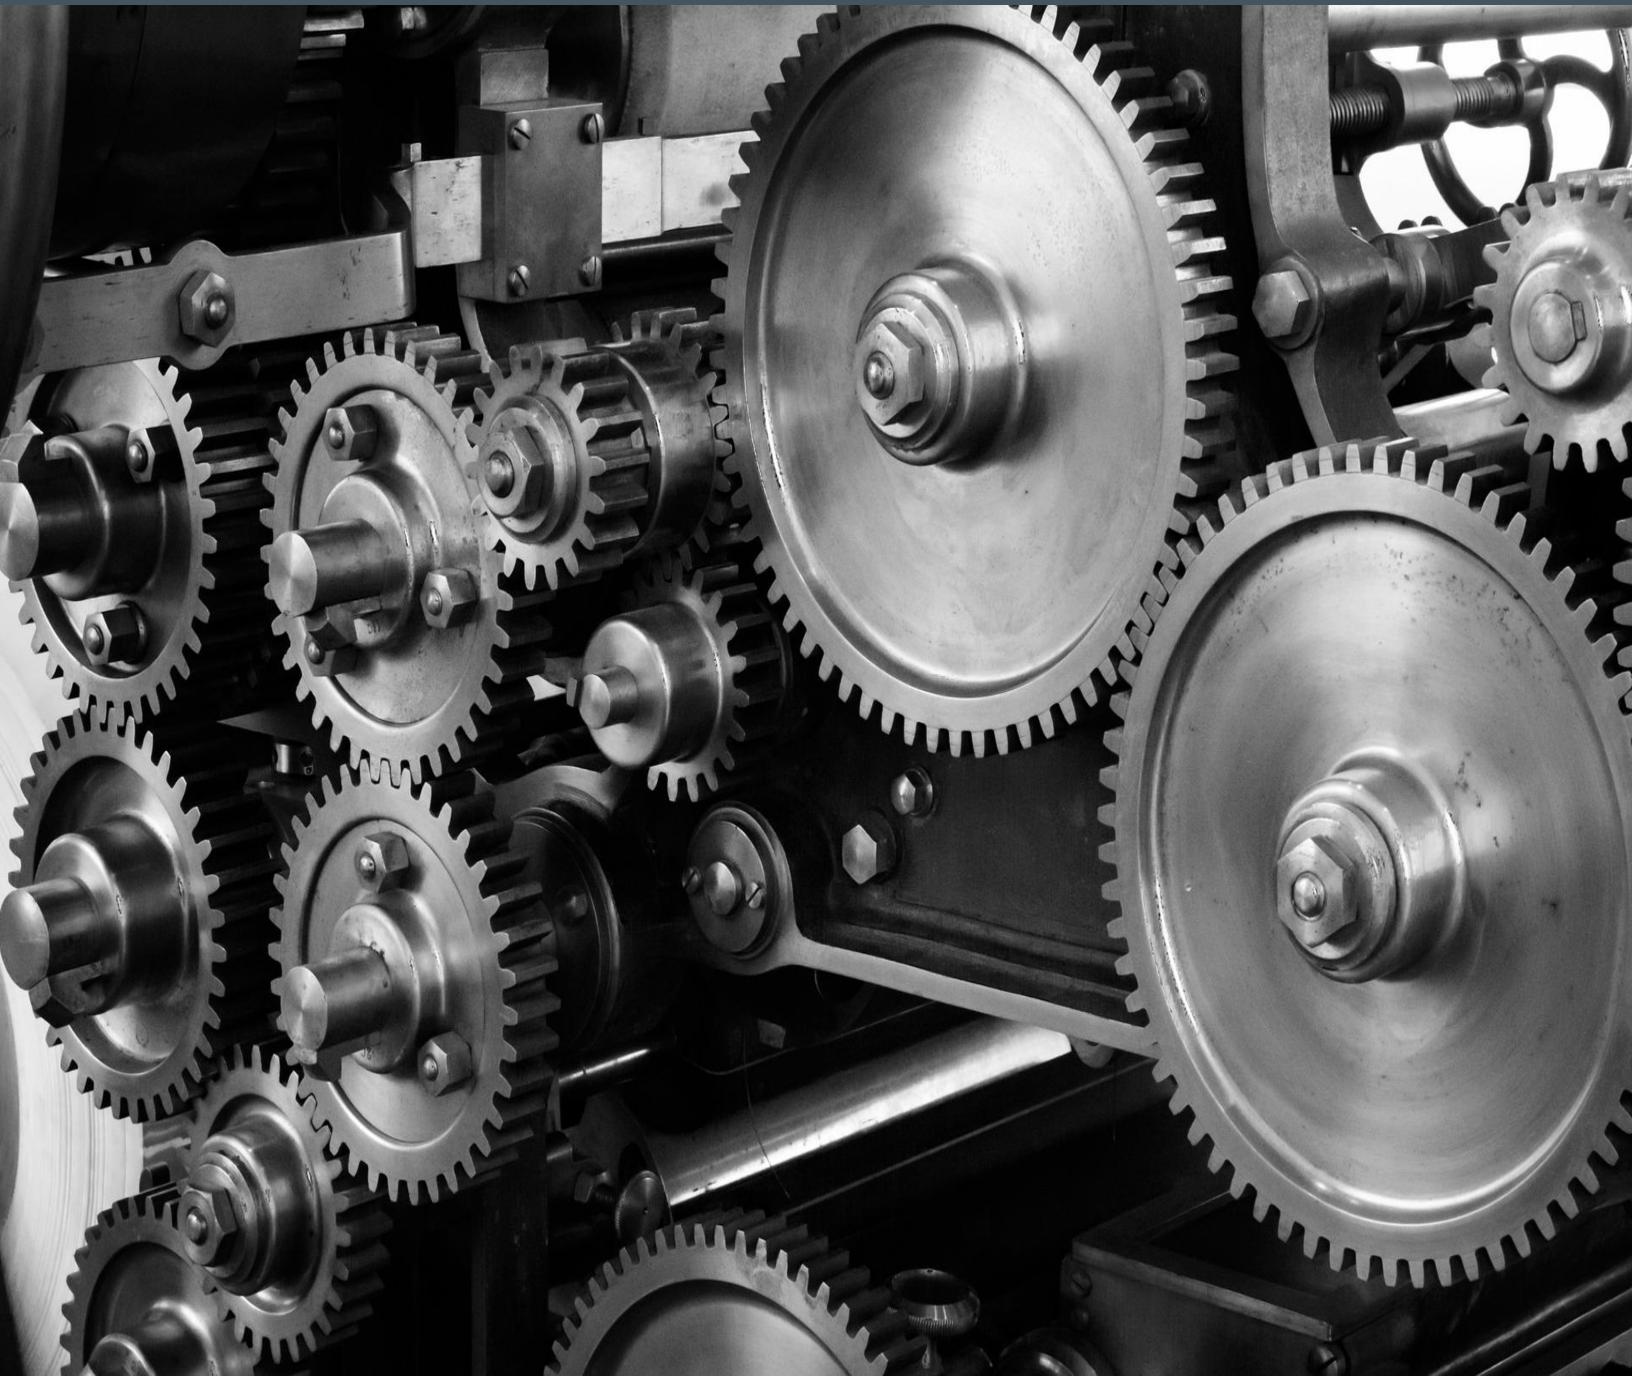

Financial Risk Group | 264 W. Chatham Street Cary, NC 27511 | 919 439 3819

Machinist (noun): A person who operates a machine, especially a machine tool

As a journal that seeks to provide insight into the latest Machine Learning tools and applications available to practitioners, a primary question to address is: "How can I apply this tool set?" Or put another way, how might these tools be used in a functional process?

In this issue of the New Machinist Journal, we explore a unique application of machine learning (ML) within the financial services industry. This industry has been open to the concept of machine learning, and a voracious consumer of the technology in investment and trading, in particular, but its users have been slower to adopt ML applications (outside of those that provide immediate feedback by way of profit and loss). Let's examine how ML could be applied to Trade Desk Surveillance, an area of risk management.

Risk management mentions in the popular press unfortunately often come with a negative preface— e.g., "a lapse of risk management," or "inadequate risk management." Machine Learning can help mitigate these tags, by providing new perspective and bringing to light aberrant behavior. There are a number of areas that benefit from this type of application, but one, in particular is a natural fit: the trading desk.

Many of the functions that traders once facilitated have been supplanted by technology. The niche traders that remain have seen their influence grow, primarily *because of* technology. Though their expanded influence can have positive effects, from a risk management perspective we should consider the negative implications. For example, traders are often monitored for a number of activities that might lead to

- Exceeding Product Limits
- Exceeding Risk Limits
- Front Running
- Wash Trades
- Offsetting Trades
- Trade Concentration
- Marking the Close
- Out-of-Market Transactions
- Percent-of-Market Activity
- Open Interest Violations
- Cancellations, Fictitious Orders, Spoofing
- Price Ramping
- Market Movements
- Modified/Late/Missing Trades

regulatory and reputational issues, including:

These are currently surveilled using "rules engines" with defined parameters that, if exceeded, alert management to the need for additional investigation. But there are two primary issues with this approach. One, the engines utilize static rules that are only updated periodically (often annually). Two, these engines often generate alerts at day's end or even next day. Machine Learning can be employed to address both of the shortcomings.

The static rules engine is ripe for augmentation through the use of ML. Rules could be automatically adjusted to reflect changing macro environmental factors, as well as the micro economic factors of a trader's individual behaviors. To better understand the application there is a need to explain the problem of static rules in more detail.

Rules engines typically apply to macro parameters such as per-trade or per-day trading limits, either in aggregate or a net basis. However, this provides very little insight into the behavior beyond the obvious, which an unethical trader with minimal experience will learn how to skirt. Utilizing ML can detect behavioral aberrations that stay within the limits—but deviate from the expected action. These deviations may be point, contextual, or collective anomalies, depending on the specific area of interest.

Consider this simple example: a fixed income trader who normally trades 80% of the time at the far end of the curve suddenly changes to the short end, but still trades within the established boundaries. A static rule wouldn't detect an issue. But a rule that *learned* the trader's behavior would detect this anomaly and alert the appropriate parties. Machine Learning allows for these changes in behavior to reflect the current environment, keeping false positives to a minimum. A Support Vector Machine may provide the context necessary to identify those instances which fall outside of the expected cluster in such a dynamic setting.

Regarding the second point of timeliness: ML is facilitated through dynamic information, and as such can process and identify probable issues in real time. In the case of a trading desk this is the difference between addressing an issue during the trading day, versus post trading or even next day. It can help catch inappropriate acts sooner and reduce the possibility of reputation risk.

These very simple examples of ML risk management application in a financial services process provide some context for the rich use of this technology. Within this industry alone there are a myriad of other applications in the front office, as well as the middle and back office areas. The challenge is in understanding the processes, the problems, and the improvement opportunities, and then approaching them with the new machinist's tool sets.

Future editions will discuss ML in other areas of risk management. Stay tuned...

For more information on FRG's research in the areas of Artificial Intelligence and Machine Learning, please visit the website or contact the FRG Research Institute at Research@frgrisk.com.