# BERKERYNOYES

# How Julia Language Is Altering the FinTech Market

## White Paper

## Introduction

Julia Language promises to speed up numerical computing in the financial markets, spurring innovation in risk management and algorithmic trading.

Every year new technologies emerge, and some of them completely disrupt the marketplace. In the case of new software programming languages, the potential disruption of a given market must wait until enough developers adopt it. That tipping point may have arrived as Julia, a new open source programming language created at MIT Computer Science and Artificial Intelligence Laboratory, gains a critical mass of adherents. The first compiler was released in 2012 and currently ranks #51 on the RedMonk language ratings. Julia's key feature is that it is a high-level dynamic code that can outperform other similar JIT (Just In Time) interpretive languages.

## Faster and Easier

The innovation of Julia is that it addresses the need to easily create new numerical algorithms while still executing fast. Julia's creators noted that, before Julia, programmers would typically develop their algorithms in MATLAB, R or Python, and then re-code the algorithms into C or FORTRAN for production speed. Obviously, this slows the speed of developing usable new algorithms for numerical applications. In testing of seven basic algorithms, Julia is impressively 20 times faster than Python, 100 times faster than R, 93 times faster than MATLAB, and 1.5 times faster than FORTRAN. Julia puts high-performance computing into the hands of financial quants and scientists, and frees them from having to know the intricacies of high-speed computer science.

So how can FinTech use such a powerful technology?

## Financial Models and Regulatory Compliance

Last year at the ISC High Performance Computing event, Erik Vynckier, Chief Investment Officer (Insurance) of AllianceBernstein, presented his thoughts on the use of high-performance computing and Big Data for risk analysis to help financial institutions meet new regulatory requirements. This buy-side risk analysis includes the ability to price complex risks on the fly (collateral management, CVA) and dynamically hedge multi-asset risks (variable annuities). He posited that Big Data and high-

performance computing are critical to capital modelling, which requires Monte Carlo simulations of the complete balance sheet covering investments and liabilities at a granular level. In addition, he foresees the use of Big Data techniques in algorithmic trading and investing, and in analyzing contagion across financial networks. The financial industry is tackling these problems with advanced mathematics and increased computational capacity.

Another big challenge for FinTech is providing technology to help comply with new regulations implemented after the 2008 financial crisis. In December, the Swift Institute published 'The Role of Big Data in Governance: a regulatory and legal perspective of analytics in global financial services" by Dr. Daniel Gozman, Professor Wendy Currie, and Dr. Jonathan Seddon, which discusses the use of Big Data for compliance. In the paper, the authors discuss how Charles River Development has created a Compliance as a Service System based on Big Data technologies, which leverages large amount of data collected from their OMS. Charles River's product allows for clients "to run what-if scenarios on historical data, complex derivative analysis and check compliance rules at any stage of the trading process." The report also states:

*"One of the strengths of the CRD compliance tool is the ability for customers to write their own tests with an English-like syntax developed by the company. This, combined with the library of 1,700+ templates for 35 regulatory bodies (including rules from regulations such as UCITS, MiFID, Dodd-Frank, and Sarbanes-Oxley) meant that each customer could quickly trade new types of security and comply with differing requirements. These compliance checks can now be run at any stage of the trading cycle and can even be used to take a holding in a particular security to the limit of what was allowed. There are always rules that are subject to a company's own interpretation (such as the fair distribution of a partial executed order) however a full audit trail of every change to the order during its life is also recorded in the system."*

Clearly, Big Data has a role to play in the capital markets, and the ability to quickly modify high-performance numerical algorithms is very important. Julia is well-positioned to be the language of choice for these tasks.

**NY Fed Turns to Julia**

One of the first noteworthy applications of Julia with financial Big Data occurred at The Federal Reserve Bank of New York. In December 2015, the bank stated that it is now using Julia to create its DSGE model (Dynamic stochastic general equilibrium). Fed programmers have converted this model from MATLAB, resulting in calculations that the bank says are 10 times faster.

Another FinTech innovation hotbed is algorithmic trading (aka high-frequency trading), a technique that has seen explosive growth in the past 15 years, reaching 60 percent of all U.S. equity volume in 2010 before retreating to about 50 percent after the so-called Flash Crash. This form of trading depends on high-speed numerical computing using rapidly evolving quantitative finance algorithms which take advantage of the smallest changes in market conditions. Could Julia help speed the development and deployment of new algorithms that can outsmart other players in the market? I asked Viral B. Shah, one of the creators of Julia, if Julia is applicable to high-frequency trading. It is, he told me, adding that Julia still needs more development of memory management.

**Julia Beyond FinTech**

Outside of FinTech, Julia is already in commercial use, most notably by Soasta, a Silicon Valley company with over $74 million in equity investment. It created a Big Data product for web and application analysis and is currently using Julia for its "Data Science Workbench," which is used to analyze billions of records of beacon data generated by web and mobile users. Phillip Tellis chose Julia over R and MATLAB as its platform, due to its superior performance.

Intel is providing support for Julia to operate on its High Performance Computing chips by incorporating it into their Math Kernel Library (MKL). This platform helps take advantage of Julia's parallel computing features. Just this March, the high performance scripting team at Intel Labs released ParallelAccelerator.jl, a Julia package for parallel high-performance, high-level array style programming. Intel chose as the first use case of this library, the Black Scholes option pricing Benchmark, a Noble prize winning mathematical model of derivative financial instruments.

To aid the adoption of Julia, its creators--Viral Shah, Jeff Bezanson, Stefan Karpinski, and MIT applied mathematics

professor, Alan Edelman--have launched a startup called Julia Computing, Inc., which has presence in Boston, New York, Bangalore, San Fransisco and London. Julia Computing has products for seamless development, deployment and support of Julia based software. Julia Computing also provides a variety of professional services around Julia. In November, it received a grant from the Gordon and Betty Moore Foundation, established by the co-founder of Intel and author of "Moore's Law" of computing development.

**FinTech Marketplace**

Other companies positioned to take advantage of Julia for FinTech include 1010Data, whose platform is used to monitor financial risk; OneTick Data, which provides data for quantitative research into the financial markets and currently integrates with R and MATLAB; and Tibco, which analyzes streams of capital markets data. In addition, major FinTech players could take add Julia to their services including: Bloomberg, Thomson Reuters, Morningstar, MSCI, Interactive Data Corporation, Envestnet, FIS, McGraw-Hill Financial, Moody's, Verisk, FactSet, and FinInfo.

**Looking Ahead**

Currently, the Julia ecosystem is still immature, but numerous programmers are developing new Julia packages for the financial markets. Modern applications also need to be deployed in the cloud, and to be work as distributed virtualized applications. Cloud features are built into Julia, and Julia applications are easily integrated with cloud containers (including Docker). Another requirement for modern languages is to support Service Oriented Architectures (SOA), including micro-service architectures. Julia Webstack provides a platform for developing SOA and micro services in Julia. In addition, the developer community is growing. At this year's third JuliaCon conference at MIT, programmers will gather for technical talks, workshops and to share experiences with this new language.

As promising as Julia may be, other new or improved technologies may challenge Julia for FinTech supremacy. Faster hardware speed may make R, MATLAB and Python competitive with Julia. According to Dr. Jonathan Seddon of Audencia Business School, a large amount of high-

speed trading development is focused on hardware with the CPU being replaced with the GPU and FPGA which would enable parallel processing. He also states that hardware improvements with non-volatile memory and massively scalable parallel file systems will also speed the use of Big Data. Can Julia take advantage of these hardware changes? We know that Julia has been architected for parallel processing. Can the other languages match Julia? Dr. Seddon also mentions the development of FinQL as a computer language specifically created for financial data, which may offer a serious challenge to Julia.

### Conclusion

Despite the possible alternatives on the horizon, the future of Julia appears to be bright. It is a technology poised to take advantage of new hardware innovations and put that power in the hands of financial data scientists and entrepreneurs. It appears to be well-positioned to disrupt the current FinTech market and help financial institutions be innovative in the use of powerful new algorithms to leverage financial Big Data.

### About Berkery Noyes

Founded in 1980, Berkery Noyes is an independent investment bank that provides M&A advisory and financial consulting services to middle market companies in the information and technology industries. The firm offers skilled transaction management to publicly traded and privately held businesses and private equity groups in both sell-side and buy-side transactions. Berkery Noyes has managed over 500 transactions, ranging from several million to more than four billion dollars in value.

### About Keith Lubell

Keith Lubell is Chief Technology Officer at Berkery Noyes. Since 2002, Keith has been developing an application and database of the Mergers and Acquisitions marketplace, encompassing more than 47,000 transactions and 123,000 companies in the software and information Industry, including a toolset to understand business models and identify short and long-term trends. Berkery Noyes uses this application to generate market reports and to host live access to the data via mandasoft.com. Keith earned a B.A. degree in Physics from Middlebury College.

### About Peter Ognibene

Peter Ognibene is a managing director in the Financial Technology Group at Berkery Noyes. He has managed more than 100 sell-side and buy-side transactions during his career. Peter holds a BA from Syracuse University and an MBA from New York University.

## SELECTED FINANCIAL TECHNOLOGY & INFORMATION TRANSACTIONS



4