Accepted Manuscript

Title: Some Financial Regulatory Implications of Artificial Intelligence

Author: Larry D. Wall

PII: S0148-6195(17)30261-8
DOI: https://doi.org/10.1016/j.jeconbus.2018.05.003
Reference: JEB 5811

To appear in: *Journal of Economics and Business*

Received date: 12-11-2017
Revised date: 21-5-2018
Accepted date: 30-5-2018

Please cite this article as: Wall LD, Some Financial Regulatory Implications of Artificial Intelligence, *Journal of Economics and Business* (2018), https://doi.org/10.1016/j.jeconbus.2018.05.003

Some Financial Regulatory Implications of Artificial Intelligence

I am the sole author of the article.

My name and affiliation are:

Larry D. Wall
Federal Reserve Bank of Atlanta

Alternatively, The Evise guide for authors requests that each author clearly indicate their given name and their family name. Along with my given first name, I also use my middle initial. Thus, the same information split

- **Given name and middle initial**: Larry D.
- **Family name**: Wall
- **Affiliation**: Federal Reserve Bank of Atlanta

Highlights

- Artificial intelligence and machine learning are transforming the financial services industry.
- Machine learning is being applied by banks to serve customers and comply with regulations.
- Government agencies are using machine learning for supervisory purposes.
- The application of artificial intelligence to regulation will be more limited for the foreseeable future.
- Artificial intelligence also has the potential to substantially change both the customer base of financial firms as well as the structure of the financial industry.

1. Introduction

Andrew Ng recently compared the transformative power of artificial intelligence (AI) to that of electricity saying "Just as electricity transformed almost everything 100 years ago, today I actually have a hard time thinking of an industry that I don't think AI will transform in the next several years."[1] Although Ng's timescale may be a bit optimistic, those who have studied recent developments in AI generally agree that it will have a transformative effect on a wide variety of industries. Techniques developed in machine learning (ML), a subfield of AI, recently achieved considerable public attention with their success in playing the games Go (Hassabis, 2016) and Poker (Condliffe, 2017). Moreover, ML is being used in large-scale production processes such as Amazon's voice recognition, Google's search engines and Netflix's movie recommendations.

One industry that has attracted considerable interest and is in the early stages of being transformed is the financial services industry. Indeed, Economist (2017) recently proclaimed, "Machine-learning promises to shake up large swathes of finance." The transformations induced by AI and especially ML are also likely to have implications for financial supervisors concerned about the conduct and/or the prudent operation of financial firms. At a minimum, supervisors will need to take account of the opportunities for enhanced compliance and safety created by AI, as well as be aware of the ways that AI could be used undermine the goals of existing regulation. However, ideally the development of AI will do more than just challenge the supervisors to keep up with industry; it will also create opportunities for supervisors to more efficiently and effectively deploy their resources to accomplish their missions.

This study discusses some of the issues raised by AI for prudential supervisors with a focus on its most popular subfield, ML, and its subfield of deep learning. The first section provides a

[1] See Lynch (2017).

high-level overview of the current state of AI, including developments in its subfields of ML and deep learning. The second section discusses some of the ways in which AI is being applied by the financial services industry with an emphasis on applications that are relevant to supervisory concerns. This includes the use of AI to help firms comply with existing regulation—a part of a broader development often called RegTech. The third section considers the usage of AI by prudential supervisors, with an emphasis on how AI can and cannot be helpful. The fourth and fifth sections provides some speculative thoughts on how developments in AI might over the longer run change the financial services environment and broader economy.

2. Artificial intelligence, machine learning, and deep learning

The terms AI, ML and deep learning are increasingly showing up the general press. However, in order to understand how these technologies may influence financial supervision, first it is important to have at least a high-level understanding of what these technologies are, how they engage in learning, and what are some of their important strengths and weaknesses.

2.1 What is meant by artificial intelligence, machine learning, and deep learning

The terms "artificial intelligence," "machine learning," and "deep learning" have been given a variety of definitions over time and the disagreements in the literature remain unresolved.[2] For the purposes of this paper, AI will be defined as in the Oxford Dictionary as "The theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between

[2] Artificial intelligence also can be divided into narrow (or weak) AI and strong AI (or artificial general intelligence, AGI). As with other terms associated with AI, there is not a single universally accepted definition of these terms. The Future of Life Institute defines narrow AI as AI that "is designed to perform a narrow task" but that "AGI would outperform humans at nearly every cognitive task." This study focuses exclusively on currently available techniques, which is to say all of the discussion relates to narrow AI. If AGI is developed, it will have a far more profound impact on the financial system and human society than the technologies discussed in this paper. The Future of Life Institute definitions are available at https://futureoflife.org/background/benefits-risks-of-artificial-intelligence/.

languages."[3] One way of implementing AI is to develop an "expert system." That is, to build a database of knowledge from human experts and apply this data to offer advice or make decisions. This technique was popular in the 1980s but it has attracted relatively less attention as people working on expert systems have come to understand better the complexity of many seemingly simple problems.

An alternative way of implementing AI is to have the machine learn directly from the data. In 1959, one of the early pioneers in the field, Arthur Samuel, defined machine learning as the "field of study that gives computers the ability to learn without being explicitly programmed".[4] ML has come to dominate most areas of the field in recent years as improvements in computing speed, data availability and analysis techniques have facilitated greater accuracy at lower cost.

ML encompasses a variety of methods that can be employed for a variety of purposes including predicting values, classifying objects, discovering structure, and finding unusual data points. Even within each of these purposes, there are a variety of techniques that can be used depending upon available resources and the exact problem. Some of these techniques, such as linear and logistic regressions, would be familiar to anyone who has studied basic statistics. However, most economists are not as familiar with some other techniques, such as CART and random forest.

A potential problem with many machine-learning techniques, such as regression analysis, is that these techniques make assumptions about the structure of the data being analyzed. These assumptions may not hold even approximately for some problems and may even be irrelevant for

[3] See https://en.oxforddictionaries.com/definition/artificial_intelligence. An alternative definition from one of the original pioneers of AI, John McCarthy in 2007, is that artificial intelligence is: "the science and engineering of making intelligent machines…". He followed up this by defining intelligence as "Intelligence is the computational part of the ability to achieve goals in the world…". (See http://www-formal.stanford.edu/jmc/whatisai/node1.html) I prefer the Oxford Dictionary definition because of its relative clarity.

[4] Arthur Samuel is widely attributed to having given this definition in 1959. For example, see https://www.coursera.org/learn/machine-learning/lecture/Ujm7v/what-is-machine-learning.

other problems (such as determining whether a picture is a cat). To overcome these difficulties, computer scientists have developed methods loosely based on the working of the human brain to allow machines to learn for themselves. These methods are called deep learning, which this paper will define as "a subfield of machine learning concerned with algorithms inspired by the structure and function of the brain called artificial neural networks."[5]

The exact way in which neural networks work varies across different types of neural networks and is evolving through time. A simplified way of viewing artificial neural networks is that they may contain a large number of nodes (thousands or even millions).[6] Each node takes in data, and assigns a weight to individual data items. The weighted data values are summed and compared with a threshold value. If the sum is less than the threshold then the node passes through no data, whereas if the threshold is exceeded the node "fires" and passes along some value (typically the weighted sum of the inputs). The individual nodes are arranged in layers with the first layer taking in the raw data, processing it and passing the results onto higher-level nodes that then perform similar processes. The neural network typically learns by training on real data in which the correct answer is already known. The training process consists of adjusting the weights and thresholds to improve the accuracy of the forecast.

2.2 Different ways in which machines can learn

Machines can learn from the data in a variety of different ways including supervised learning, unsupervised learning and reinforcement learning. In supervised learning, the training data is labeled so that the machine learns to use the input data to predict the desired output value. For example, the machine is given a large number of pictures that are labeled as having a cat or as not having a cat. The machine uses the pictures and learns how to identify the features associated

[5] This definition comes from Brownlee (2016).
[6] This example is obtained from Hardesty (2017).

with cats.[7] Supervised learning can be used to categorize items (whether something is a cat) and to predict numerical values (such as stock returns).

Alternatively, the machine can engage in unsupervised learning in which the input data are not labeled. Probably the most common method of unsupervised learning is that of clustering, that is for finding patterns in the data.

A third type of learning is called reinforcement learning. Here the algorithm makes a decision at each step and is "rewarded" for taking good actions.[8] One area where reinforcement learning has been applied is the area of playing games such as Go and Poker.

2.3 Some of the strengths and limitations of machine learning

In order for most contemporary economists to understand the strengths and weaknesses of ML, it is helpful to compare ML to statistics, a discipline better known to most economists. Wasserman (2012) raises the question of the differences between ML and statistics. His short answer is "None. They are both concerned with the same question: how do we learn from data?" He proceeds to argue that there are difference between the ML and statistics in practice, but that these differences are more due to historical and sociological reasons. Statistics arose before computers and tends to emphasize "formal statistical inference (confidence intervals, hypothesis tests, optimal estimators) in low dimensional problems." In contrast, ML originated in computer science departments and "emphasizes high dimensional prediction problems." However, he observes that the differences between the fields become "blurrier all the time."

[7] Mordvintsev, Olah, and Tyka (2015) provide a high-level description of how a neural network identifies what is in a picture. They note that each layer in the extracts higher and higher order features of the image. The first layer may look for edges or corners, the intermediate layers interpret the features to look for overall shapes such as that of a door or leaf. The last few layers assemble the information on these shapes to provide complete interpretations.

[8] One could think of rewards as positive or negative points that are assigned to the transition to certain states. The algorithm controlling the weights and thresholds for individual neurons seeks to maximize these rewards (maximize an objective function based on the rewards), allowing it to learn how to better achieve whatever goal it is being assigned.

2.3.1 Machine learning versus statistics

The difference between statistics' emphasis on hypothesis testing and ML's emphasis on prediction lead to both a strength and weakness in ML relative to statistics. Historically economists have developed empirically testable hypotheses from their theories and used statistics to test these hypotheses. This approach is used to overcome a well-known limitation of statistics that significant correlation between two variables need not indicate that a causal relationship exists. Empirical tests of theoretically developed hypotheses overcame this limitation by putting the burden of identifying causal relationships on theory. However, ML's emphasis on prediction problems has led it to take an atheoretical approach to its analysis. As a result, ML can identify relationships that have not (yet) been identified by theory but it does so at the cost of potentially identifying relationships that are not causal and, thus, cannot be directly exploited. One way that users of ML techniques address this problem is to take an "iterative and experimental" approach in which small-scale, real world tests are run to determine which correlations identified by ML techniques can be usefully exploited in a causal manner.

Along with the limitations ML imports from statistics, ML that relies on the currently popular deep learning techniques has one other important limitation. The process by which deep learning techniques reach decisions is unclear. Deep learning techniques provide predictions but they do not provide insight into how the variables are being used to reach those predictions (Knight, 2017a).

2.3.2 Machine learning versus human capabilities

Given that ML is a form of statistics, some well-known strengths and limitations relative to human capabilities follow almost immediately. One strength of ML and statistics are that they are designed to process far more data than could be done by any human. This allows machine-learning techniques to identify empirical relationships that humans could easily miss. A second advantage is that the computer provides a level of consistency that is not possible for humans who are sometimes distracted by hunger, lack of sleep or non-work related issues in their lives. A third strength of ML in large-scale operations is that it typically has substantially lower cost marginal costs than reliance on humans.

However, ML also imports the same limitations as statistics. Two important limitations for the use of ML in finance follow from Rowe's (2013) observation that "No amount of complex mathematical/statistical analysis can possibly squeeze more information from a data set than it contains initially."[9] The first of these limitations is that there need to be sufficient historical examples of the phenomenon for the empirical analysis to identify the factors that reliably predict its occurrence. As an example of a potentially important event for which we (fortunately) have too few examples, one may want to predict the probability that a developed country would default on its sovereign debt in the next five years. However, the problem of an absolute lack of events is relatively easy to identify compared with the problem of a lack of events due to a particular cause. For example, one might apply ML to a bank's historical credit experience to help predict delinquencies and the model might prove highly accurate given these data. However, before a loan could show up in the bank's database, a customer must first have applied for the loan and the bank must have granted the loan based on various known and possibly unknown criteria. Thus, the data may not provide very good estimates for customer groups that would not historically have applied to that bank for the loan or who the bank had historically screened out prior to the completion of an application.

[9] Reinforcement learning can avoid this problem by generating its own data, such as having a machine compete against another machine in the game of Go.

The second limitation that follows from Rowe is that for ML to predict a phenomenon, that phenomenon must be labeled in the data. This can be a problem, for example, if one knows that a data set of transactions contains a significant number of cases involving fraud but cannot say which transactions were fraudulent. This limitation can be somewhat mitigated in some cases, however, by the use of unsupervised learning. For example, unsupervised ML techniques may be applied to a dataset to identify a set of transactions with characteristics that are different from the others. Unusual transaction may then be further analyzed to determine whether they are a result of fraud.

3. Applications of machine learning by banks

Although the application of ML to financial problems is relatively new, banks and other financial firms have begun exploring and using ML in a variety of ways. The following subsections highlight some of the ways banks are or could use ML to serve their customers better and to meet increasingly demanding regulatory requirements.

3.1 Some uses of machine learning to serve customers

Banks are using ML in a variety of ways to serve their customers. Some of these applications have been in place for years. For example, van Liebergen (2017) reports that banks have been using ML techniques for over a decade to detect credit card fraud with some significant success. Some other relatively new uses are similar to those being employed by nonfinancial firms, such as Bank of America's development of the chatbot "erica."[10] However, ML has potential applications in a variety of areas in the financial services that raise unique regulatory questions.

One obvious area that could potentially benefit from ML is that of measuring credit risk.[11] A lender that could identify customers that are currently paying higher credit risk premiums than is justified could gain profitable market share by offering these customers lower price loans.

[10] See Crosman (2017) for a discussion of erica.
[11] See Jagtiani and Lemieux (2018) for an analysis of the use of big data and machine learning by the LendingClub.

Similarly, a lender that could identify customers that are being undercharged for credit risk could reduce its losses by charging them more or denying their loan request. However, the use of ML for lending also raises several potential problems. First, the data used to train the machine-learning algorithm may not be representative of the range of customers that will apply for the loans leading to an increased error rate. A second problem is that as people learn how the model works, the higher risk borrowers can learn to mimic the behavior of lower risk borrowers before applying for a loan. Both of these problems are relatively well understood among experienced lenders and can be at least somewhat mitigated by careful monitoring of delinquencies after the loans have been granted.

Lending decisions might be significantly improved by the application of current deep learning techniques; however, the lack of transparency in these models is a potentially large obstacle. The U.S. prohibits discrimination based on various categories including race, sex and marital status. Moreover, a lending algorithm could be found in violation of this prohibition even if the algorithm does not directly use any of the prohibited categories but rather uses data that may be highly correlated with protected categories, such as grammatical errors in the lending application according to Petrasic et al. (2017).[12] The lack of transparency could become an even more difficult problem in the European Union (EU) where the General Data Protection Regulation took effect in 2018 gives their citizens the right to receive an explanation for decisions based solely on automated processing according to Goodman and Flaxman (2016). Various efforts are underway to mitigate the lack of transparency and make deep learning results more transparent.[13]

[12] The U.S. Consumer Financial Protection Bureau (2017a) issued a request for information on the use of alternative data and modeling techniques (including machine learning). A no action letter was issued by the U.S. Consumer Financial Protection Bureau (2017b) with regards to the use of alternative data and new technologies such as machine learning by Upstart Network. The no action letter comes with requirements for monthly reporting to the agency.

[13] See Knight (2017b) for a discussion of some efforts to make deep learning models decisions more transparent.

Another use of ML is to develop strategies for investing and order execution. Kirilenko and Lo (2013) discuss the developments in financial and computing technology that have led algorithmic trading to become a major part of trading in the financial systems. Initially these computer models relied on human programming. However, ML is coming to play an increasingly large role. For example, the world's largest asset manager, Blackrock, announced that it is going to rely more on computers to pick stocks and that it was laying off 40 employees including portfolio managers.[14] More recently Noonan (2017) reports that JP Morgan will use ML to execute trades for its customers in equity markets.

The increased use of machines in investment advising and trading comes with some risks. One concern is that the application of ML could facilitate more trading errors. Kirilenko and Lo (2013) quote the technology-specific corollary of Murphy's Law: "whatever can go wrong will go wrong faster and bigger when computers are involved." They provide several examples of things going wrong, albeit none directly attributable machine learning. These examples included flash crashes in which prices suddenly and very dramatically spiked up or down for no apparent reason.[15] A deeper concern expressed by Carney (2017b) is that it could lead to "excess volatility or increase pro-cyclicality as a result of herding." His concern is that the underlying algorithms could be too sensitive to price moves or that the algorithms may produce highly correlated recommendations.

ML is also being used in RegTech (regulatory technology) to reduce the cost and increase the effectiveness of compliance with various regulatory requirements. One regulatory area that

[14] See Reuters (2017) for more on the announcement and Segal (2016) for a discussion of Blackrock's use of AI and ML.

[15] Another example of a problem made worse by automation given by Kirilenko and Lo (2013) relates to Knight Capital Group, Inc. on August 1, 2012. Knight encountered a problem at the opening of the market with some software installation that resulted in a large volume of unintended trades. Knight was not able to void the trades, so it had to liquidate its positions. The resulting losses wiped out most of Knight's capital and forced it to sell itself to GETCO.

has been especially challenging for many large U.S. banks is that of the Comprehensive Capital Analysis and Review (CCAR). CCAR not only requires banks to demonstrate they would remain adequately capitalized through a stressful scenario; the banks also have to demonstrate that they have a “robust forward-looking capital-planning process” which requires adequately documenting the processes used in their modeling.[16] ML techniques helped Citigroup pass CCAR by improving the way the bank developed its internal models.[17]

Another area where ML is being applied is in the area of conduct and market abuse in trading. Carney (2017a) states that global banks misconduct costs have exceeded $230 billion. One of the responses of banks according to van Liebergen (2017) is to develop automate systems that monitor a variety of behavior by traders. The behaviors may include trading patterns, e-mail traffic, calendar items, and even telephone calls. Among the challenges discussed by van Liebergen (2017) is that often there is a lack of labeled data for supervised learning and the need to be able to audit the results.

4. Financial supervision and regulation

The terms “supervision” and “regulation” are often used interchangeably as a shorthand for “supervision and regulation.” However, these terms relate to different activities and those differences have important implications for the potential contribution of ML.

Regulation is a formal process of writing the rules that define acceptable behavior. In the U.S., the Congress often provides the federal financial agencies with a general set of goals and specific tools to obtain those goals. The agencies then write relatively more detailed sets of

[16] Quote taken from https://www.federalreserve.gov/supervisionreg/stress-tests-capital-planning.htm.
[17] See Arnold (2017). See also Woods (2015) for a discussion of how ML can be applied to revenue forecasting models for CCAR.

requirements that specify a range of conduct that is not acceptable behavior (or alternatively, a range of behavior that is acceptable).

The enforcement of regulations is referred to as supervision. The federal financial agencies typically use some combination of off-site data analysis and on-site examinations to evaluate compliance with the regulations. Whether a financial firm complies with a regulation is sometimes readily determined, especially if the regulation draws a bright boundary between acceptable and unacceptable behavior based on readily observable facts. However, the determination of compliance with other regulations is not so straightforward. For example, a bank may report that it complies with minimum regulatory capital requirements but examiners will want to check periodically to confirm that the bank has not overstated its capital by delaying recognition of some of the deterioration in its asset values. Moreover, in some cases the boundary between acceptable and unacceptable is rather fuzzy, as for example the requirement that a bank operate in a "safe and sound" manner. In these cases, the agency may readily identify potential problems but need to engage in further examination and discussion with the bank to determine if this requirement is being violated.

Given these important differences, the following two subsections discuss the use of ML in supervision and then in regulation.

4.1 Supervision

ML can be a valuable input into financial supervision by helping financial supervisors to identify issues that need further analysis. The financial supervisors can evaluate the issues identified by ML based on their accumulated knowledge about the relevant markets and/or institutions. Additionally, the financial supervisors may conduct further analysis using various other inputs, including discussions with industry participants, to evaluate these issues.

The potential for ML to provide a useful input is highlighted by the longstanding application and proven usefulness of conventional statistics to analyzing financial data. Various aspects of the financial system have been analyzed in the thousands of empirical papers written on this topic. Some of these papers help us to better understand the financial system but have no direct ties to financial supervision. However, other parts of the literature have contributed directly to financial supervision. For example, the large accounting literature on earnings manipulation has contributed to the ability of the Securities and Exchange Commission (SEC) to identify violations of that agency's disclosure requirements. Another example is the so-called early warning literature that seeks to identify banks that are more likely to become distressed or fail. Bank supervisors have used insights from this literature to allocate more supervisory resources to the banks most likely to benefit from those resources.

The usefulness of standard statistical analysis is limited relative to ML in at least two important ways. First, statistical analysis depends upon being able to represent the data in numerical form with categorical data typically converted to discrete numbers. ML can apply natural language processing to deal directly with words. For example, Bauguess (2017) discusses the application of ML to the SEC's file of tips, complaints and referrals. This application helps to identify common themes that allowed the individual reports to be tagged.

A second difference between ML and statistics is that almost all of the statistical analysis comes from tests of empirical hypotheses derived from theory. As such, the relationships uncovered are likely to be causal (albeit not guaranteed in all cases), but statisticians often ignore large numbers of possibly relevant variables for which the theory has no predictions. ML in contrast lets the available data speak for itself, potentially revealing important relationships that have not yet been identified by theorists. Both the unsupervised and supervised versions of ML can be useful. The unsupervised version can help by clustering observations into groups allowing further analysis of the individual groups and of the outliers that do not appear to fit into any group. The supervised version can be even more useful in helping to identify potential violations of regulations. Bauguess (2016) discusses the SEC's work to apply ML to the analysis of the narrative disclosures in financial statements to help determine the risk of violations of various disclosure rules.[18] He observes, however, that the SEC is using ML only to flag activities that might violate existing regulations and not to be a "Robocop" that automatically imposes penalties without further investigation by humans.

4.2 Regulation

ML as an atheoretic application of statistics provides some potentially useful benefits that cannot be obtained from conventional statistics. However, the atheoretical nature of ML also imposes significant limits on the use of ML in writing regulations. Additionally, ML and statistics share an important common limitation in that both depend upon available data. The next two subsections discussion the implications of ML being atheoretical and of ML being limited by available data.

4.2.1 Benefits and limits from atheoretical analysis

The potential benefit of atheoretical ML analysis arises from its imposition of less structure on the empirical analysis and thereby letting the data speak for itself. As a result, ML has the potential for uncovering previously hidden relationships that allow for a better understanding of financial markets and institutions. In supervised learning this benefit not only includes the possibility of identifying relationships with new variables, it also includes the potential for better understanding of non-linear relationships (including threshold effects) and uncovering previously

[18] A possible limitation of this approach, however, is that those filing the disclosure statements will learn what sort of language is likely to trigger

unrecognized interactions among the variables. Additionally, unsupervised learning can reveal commonalities across seemingly different groups, as well as highlighting activities and firms that are outliers.

The cost of letting the data speak for itself without the constraints of theory is that the correlations it identifies need not be causal relationships. The problem with using such correlations in the writing of regulations is that regulation are intended to impose binding limits on individuals' behavior. Yet if the relationship identified by ML is not a causal relationship then the risk arises that the regulation will impose costly constraints without necessarily helping to contribute to the underlying public policy goal of the regulation. Further, once a regulation is written, it is not easily rewritten—that is regulation writing is rarely "iterative and experimental." In part, the problem with rewriting the rules arises from the normal slow pace of any bureaucracy, and in part because the rule writing process typically requires that changes to regulation be issue in a proposed format with a comment period for the public to respond. Moreover, regulations often produce clear winners that gain a competitive advantage by optimizing their operations given the constraints imposed by the regulation. These winners may fight to keep even inefficient regulations in place. Thus, regulations are rarely written with the idea that they will be revisited and likely rewritten in the near future to reflect new information generated by the regulation itself.

4.2.2 Limits imposed by available data

Rowe's (2013) admonition about the limits a data set imposes on statistics, and by implication ML, has some important consequences for ML's usefulness in helping to write regulations. ML can provide at best limited assistance on issues where we have little or no data. Unfortunately, as Wall (2016) observes, some of the most important questions in financial regulation relate to issues on which we have limited data.

Two of the primary goals of financial regulation are the preservation of financial system stability and the prevention of significant losses due to the failure of individual large financial institutions. Fortunately, from a social perspective, bouts of financial instability and large losses are rare tail events. However, that implies that almost all of the data we have are from normal times when the financial system and the large institutions are not under stress. ML can use these data from normal times to help identify those variables that are useful in predicting losses during normal periods. However, in order to use these results to reduce the risk of instability and large institution failure, we need some theory or parametrical statistical structure linking the data obtained in normal times to the determinants of large losses that could threaten overall stability and individual institutions.

The regulation of financial firms' capital is one good example where available data limits what ML can do to improve regulation. The Basel capital accords set minimum capital requirements with the goal of ensuring that equity capital levels will remain non-negative with an over 99 percent probability. The U.S. supplements the Basel standards with stress testing designed to make sure that a bank has enough capital to not only remain solvent but also continue lending even in the event the economy undergoes a severe recession. Both of these ways of measuring capital adequacy require the projection of losses in parts of the distribution where banks have relatively little data. Thus, in both cases the supervisors and banks rely on theories about the distribution of losses to link the data that is available in abundance on normal times to what we might expect in situations where we have little or no direct experience. Although ML has proven useful in analyzing data from normal times, no improvements in ML techniques or increases in data from normal times can replace the need for theories or parametric statistical structure linking what we can observe to what might happen in extreme cases.

Another important question in writing regulations is evaluating how regulatees and others will respond to changes in regulation. The goal of regulation is to change behavior by imposing binding limits on legally acceptable behavior. In response to a change in regulation, agents are likely to seek to re-optimize their behavior given the new constraints imposed by regulation. However, this re-optimization is likely to involve not only the intended changes in behavior but also unintended changes by both regulatees and others involved in related activities. These unintended changes may lead to changes in the structure of the relevant financial markets and institutions that have important implications for the effectiveness of the regulation and possibly even for the effectiveness of other regulations. Unfortunately, the data that ML uses to make its predictions at the time a new regulation is being written are necessarily drawn from a market structure optimized for the old regulations. Hence, ML cannot predict whether or how a new regulation will change the structure of financial markets or institutions. Nevertheless, ML can still be helpful in writing regulations to the extent it helps regulators better understand current behavior and this helps them predict responses to changes in regulations. Machine learning may also help regulators in identifying some of the unintended consequences leading to a faster response.

5. The importance of data

Given that ML's ability to extract insights is limited to the information contained in the dataset it is analyzing, the quantity, quality and diversity of data is an important determinant of the insights that can be obtained from machine learning. Indeed, an increasingly popular phrase that highlights the importance of data is that "data is the new oil" of the modern economy.[19]

The fintech industry in general, including ML applied to financial problems, has seen an explosion in competition. New firms are using ML to enter the financial services industry and

[19] For a discussion of the origin of this phrase, see https://www.quora.com/Who-should-get-credit-for-the-quote-data-is-the-new-oil.

existing firms are using ML to enter new subsectors of financial services. Many of these entrants will fail due to inferior business plans and/or inferior execution. However, the benefits of access to more data suggests there may also be substantial economies of scale for financial firms that rely on ML for critical tasks such as obtaining customers and managing risks. Indeed, the behavior of the tech firms that are currently leaders in ML suggests that they perceive data as an important source of competitive advantage. Simonite (2017) observed that while some of the large tech firms such as Google and Microsoft have made their software available to others, they are "hoarding" those data sets that are of the most commercial value.[20]

To the extent that the application of ML to ever larger datasets conveys substantial competitive advantages in financial services, that could have significant implications for the structure of the financial services industry. Firms that provide the best ML enabled products will be able to gradually gain market share and in doing so obtain even more data with which to improve their ML predictions and competitive position. The end game could be a tenuous existence for the smaller firms competing against the ML giants resulting in reduced competition in financial services. If ML conveys such a competitive advantage, the resulting winners could become substantially larger and their financial condition even more important to overall financial stability (i.e., it could make the too-big-to-fail problem even worse).

To the extent that data hoarding conveys a competitive advantage, one way of limiting this advantage is by reducing the extent to which individual firms have exclusive access to data.[21] The simplest approach, that of forcing everyone to share all of their data, is likely not feasible in many

[20] Some hedge funds are also seeking and obtaining exclusive access to some datasets to improve their trading performance according to Fortado, Wigglesworth and Scannell (2017).

[21] For example, Wildau (2017) reports that the People's Bank of China has ordered online payment groups to funnel their payments through a centralized clearing house. He quotes one fintech analyst as saying this will likely result in payments information being shared with commercial banks, thereby limiting the market power of the online services Alipay and Tencent.

countries (and arguably not desirable) because of its implications for customer privacy.[22] Simonite (2017) suggests one way of reducing the advantage of larger firms is for the smaller firms to pool voluntarily their data, as is sometimes done by smaller insurance firms. Such pooling could reduce the competitive advantage of larger firms but could also raise privacy considerations similar to those of forced data sharing.[23] Another alternative would be for the legal system to take the position that the customer owns their data and can share it as they chose. The European Union in its Payment Systems Directive 2 (PSD2) has recently taken this approach. The limitations of allowing customers to share their data is that the customer must find that the personal benefits they receive from sharing the data outweighs the loss in privacy and other potential risks.

6. Impact of ML through its effect on the rest of the economy

ML is likely to produce economic winners and losers in the broader nonfinancial economy as is true with the adoption of many new technologies. What is potentially different about the impact of ML relative to prior technological changes, as observed by David Ng, is its capability of touching almost every part of the economy. Moreover, Manyika's et al. (2017) study for the consulting company McKinsey estimates that just over 50 percent of the activities currently undertaken in the global economy could be replaced by automation within the next 40 years, although that estimate could be off by 20 years in either direction. Whether this vision on the part of ML advocates will be fully realized remains to be seen. However, if a substantial portion of this vision is realized, it would have large effects on the economic environment in which the financial services operates.

[22] Privacy is not only relevant to individuals but also to corporations. A corporation's competitive position could be significantly weakened if other firms could observe its financial transactions.

[23] Sharing of data by small firms could lead to herding behavior on the part of these institutions, possibly resulting a the "too-many-to-fail" situation discussed by Acharya and Yorulmazer (2007). However, this potential increase in financial fragility would be at least partially offset to the extent that increased sharing allowed these smaller firms to continue operating rather than have their business acquired by one of a handful of ML giants.

ML is already producing winners in the tech industry and firms in many other industries are at varying stages of developing ML projects intended to enhance their competitiveness. Some of these firms will prove better able to execute ML and apply it to their businesses in ways that give them significant competitive advantage. In doing so, these firms will be seeking skilled people to prepare the data, develop the models and apply the models to their businesses. Individuals involved in this process are likely to do well in the labor market. These winning firms and individuals are likely to become profitable consumers of increased financial services.

However, to the extent that ML helps determine which firms are the winners that gain market share and profitability, it will result in other firms losing market share and profitability. Some of the losers in this process are likely to face the prospect of having to sell out to firms that are more successful or be at risk of failure. Just as financial firms can profit from timely identification of the winners, they can also avoid losses by timely identification of the losers.

Similarly, although ML has the potential to create some winners in the labor market, it also has the potential to produce automation that displaces many workers. There is an ongoing debate about whether new jobs will arise for the displaced workers. However, even if new jobs arise over the long run, some workers face the prospect of the destruction of a substantial part of their human capital. Moreover, ML's ability to perform human tasks is not limited to low skilled jobs. Some tasks done by well-educated, highly skilled workers can also be done using ML. For example, JP Morgan Chase, has a new program, called COIN, for Contract Intelligence, that interprets commercial-loan agreements. Prior to the project going on-line, that task required 360,000 hours of work each year by lawyers and loan officers (Son, 2017).

7. Conclusion

The rapid development and employment of AI techniques has the potential to transform the financial services industry along with many sectors in the real economy. To the extent this potential is realized, AI will have substantial implications for financial conduct and prudential supervisors. Moreover, AI has the potential to help supervisors identify potential violations and help regulators better anticipate the impact of changes in regulation.

References

Acharya, Viral V., and Tanju Yorulmazer (2007). "Too many to fail—An analysis of time-inconsistency in bank closure policies." *Journal of Financial Intermediation* 16, no. 1: 1-31.

Arnold, Martin (2016), "Market grows for 'regtech', or AI for regulation," *Financial Times* (October 14). Available at https://www.ft.com/content/fd80ac50-7383-11e6-bf48-b372cdb1043a.

Bauguess, Scott W. (2016), "Has Big Data Made Us Lazy?", Speech given to the Midwest Region Meeting - American Accounting Association (AAA), Chicago Illinois (October 21). Available at https://www.sec.gov/news/speech/bauguess-american-accounting-association-102116.html.

Bauguess, Scott W. (2017), "The Role of Big Data, Machine Learning, and AI in Assessing Risks: a Regulatory Perspective," Speech given to OpRisk North America 2017, New York, New York, (June 21). Available https://www.sec.gov/news/speech/bauguess-big-data-ai.

Brownlee, Jason (2016), "What is Deep Learning?" Machine Learning Mastery (April 16). Available at https://machinelearningmastery.com/what-is-deep-learning/.

Carney, Mark (2017a), "Remarks at the Banking Standards Board Panel "Worthy of trust? Law, ethics and culture in banking" Speech at a Bank of England Conference (March 21). Available at http://www.bankofengland.co.uk/publications/Documents/speeches/2017/speech970.pdf

Carney, Mark (2017b), "The Promise of FinTech – Something New Under the Sun?" Speech at the Deutsche Bundesbank G20 conference on "Digitising finance, financial inclusion and financial literacy", Wiesbaden Germany. Available at http://www.bankofengland.co.uk/publications/Documents/speeches/2017/speech956.pdf.

Condliffe, Jamie (2017), "An AI Poker Bot Has Whipped the Pros" *MIT Technology Review*, (January 31). Available at https://www.technologyreview.com/s/603544/an-ai-poker-bot-has-whipped-the-pros/.

Consumer Financial Protection Bureau (2017a) "CFPB Explores Impact of Alternative Data on Credit Access for Consumers Who Are Credit Invisible," Press Release, available at https://www.consumerfinance.gov/about-us/newsroom/cfpb-explores-impact-alternative-datacredit-access-consumers-who-are-credit-invisible/.

Consumer Financial Protection Bureau (2017b) "CFPB Announces First No-Action Letter to Upstart Network," Press Release, available at https://www.consumerfinance.gov/about-us/newsroom/cfpb-announces-first-no-action-letter-upstart-network/.

Crosman, Penny (2017), "B of A gives its bot time to become a banker," *American Banker* (May 15). Available at https://www.americanbanker.com/news/b-of-a-gives-its-bot-time-to-become-a-banker.

Domingos, Pedro (2012), "A few useful things to know about machine learning." *Communications of the ACM* 55, no. 10: 78-87.

Goodman, Bryce, and Seth Flaxman (2016). "European Union regulations on algorithmic decision-making and a" right to explanation"." arXiv preprint arXiv:1606.08813 (2016). Available at https://arxiv.org/pdf/1606.08813.pdf.

Hassabis, Dennis (2016), "AlphaGo: using machine learning to master the ancient game of Go," (January 27). Available at https://www.blog.google/topics/machine-learning/alphago-machine-learning-game-go/.

Economist, (2017), "Machine-learning promises to shake up large swathes of finance", (May 25). Available at https://www.economist.com/news/finance-and-economics/21722685-fields-trading-credit-assessment-fraud-prevention-machine-learning.

Fortado, Lindsay, Robin Wigglesworth and Kara Scannell (2017), "The Big Read Hedge funds," *Financial Times* (August 28).

Hardesty, Larry (2017), "Explained: Neural networks," *MIT News* (April 14).

Jagtiani, Julapa, and Catharine Lemieux (2018). "The Roles of Alternative Data and Machine Learning in Fintech Lending: Evidence from the LendingClub Consumer Platform.". Federal Reserve Bank of Philadelphia working paper No. 18-15. Available at https://www.philadelphiafed.org/-/media/research-and-data/publications/working-papers/2018/wp18-15.pdf.

Kirilenko, Andrei A., and Andrew W. Lo (2013), "Moore's Law versus Murphy's Law: Algorithmic Trading and Its Discontents." *Journal of Economic Perspect*ives 27, no. 2: 51-72.

Knight, Will, (2017a), "The Dark Secret at the Heart of AI", MIT Technology Review (April 13). Available at https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/.

Knight, Will, (2017b), "The Financial World Wants to Open AI's Black Boxes", *MIT Technology Review* (April 13). Available at https://www.technologyreview.com/s/604122/the-financial-world-wants-to-open-ais-black-boxes/.

Lynch, Shana, (2017), "Andrew Ng: Why AI Is the New Electricity," Insights by Stanford Business (March 11). Available at https://www.gsb.stanford.edu/insights/andrew-ng-why-ai-new-electricity.

Manyika, James, Michael Chui, Mehdi Miremadi, Jacques Bughin, Katy George, Paul Willmott, and Martin Dewhurst (2017), "A Future That Works: Automation, Employment, and Productivity (January). Available at https://www.mckinsey.com/~/media/McKinsey/Global%20Themes/Digital%20Disruption/Harnessing%20automation%20for%20a%20future%20that%20works/MGI-A-future-that-works_Full-report.ashx .

Mordvintsev, Alexander, Christopher Olah and Mike Tyka (2015). "Inceptionism: Going Deeper into Neural Networks," *Google Research Blog* (June 17). Available at https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html.

Noonan, Laura (2017), "JP Morgan Develops Robot to Execute Trades," *Financial Times* (July 31).

van Liebergen, Bart (2017), "Machine Learning: A Revolution in Risk Management and Compliance?" Institute for International Finance Research Note (April 27), Available at https://www.iif.com/file/19459/download?token=1JQdTzAX.

Petrasic, Kevin, Benjamin Saul, James Greig, and Matthew Bornfreund (2017), "Algorithms and bias: What lenders need to know," White and Case. Available at https://www.whitecase.com/publications/insight/algorithms-and-bias-what-lenders-need-know.

Row, David M. (2013), "Risk Management Beyond VaR" Paper presented at the Federal Reserve Bank of Atlanta 2013 Financial Market Conference (April 10). Available at https://normaldeviate.wordpress.com/2012/06/12/statistics-versus-machine-learning-5-2/.

Reuters (2017), "BlackRock Is Cutting Jobs and Banking on Robots to Beat the Stock Market" *Fortune* (March 29). Available at http://fortune.com/2017/03/29/blackrock-robots-stock-picking/.

Segal, Julie (2016), "BlackRock Is Making Big Data Bigger" *Institutional Investor* (November 1). Available at http://www.institutionalinvestor.com/article/3598029/asset-management-fixed-income/blackrock-is-making-big-data-bigger.html#/.WahmfsiGOr1.

Simonite, Tom (2017) "AI and 'Enormous Data' Could Make Tech Giants Harder To Topple" *Wired* (February 27). Available at https://www.bloomberg.com/news/articles/2017-02-28/jpmorgan-marshals-an-army-of-developers-to-automate-high-finance.

Wall, Larry D. (2016) "Prudential Regulation, Big Data, and Machine Learning," *Notes from the Vault* (November). Available at https://www.frbatlanta.org/cenfis/publications/notesfromthevault/11-prudential-regulation-bigdata-and-machine-learning-2016-11-21.aspx.

Wasserman, Larry (2012) "Statistics Versus Machine Learning" *Normal Deviate* (June 12). Available at https://normaldeviate.wordpress.com/2012/06/12/statistics-versus-machine-learning-5-2/.

Wildau, Gabriel (2017), "China targets mobile payments oligopoly with clearing mandate" Financial Times (August 9). Available at https://www.ft.com/content/3bcb5150-7cce-11e7-9108-edda0bcbc928.

Woods, Michael, (2015), "Creating Effective Revenue Forecast Models for CCAR," (March 25). Available at https://s3.amazonaws.com/cdn.ayasdi.com/wp-content/uploads/2017/02/13111852/creating-effective-revenue-forecast-models-for-ccar.pdf.