COMPETITION AND CONSUMER PROTECTION IN THE 21ST CENTURY

SAS INSTITUTE'S

RESPONSE TO QUESTION 9: CONSUMER WELFARE IMPLICATIONS ASSOCIATED WITH THE USE OF ALGORITHMIC DECISION TOOLS, AI, AND PREDICTIVE ANALYTICS

Thank you for the opportunity to submit comments on the topic of artificial intelligence, algorithmic tools, and predictive analytics (collectively, "AI"). As discussed further herein, we make several points:

1. There is tremendous consumer welfare created by the application of analytics and artificial intelligence (AI) to consumer advertising and marketing campaigns, as well as the use by firms to make pricing and output decisions. At the same time, because many of the applications do not need personally identifiable information to help make decisions that are beneficial to consumer welfare, these application, as such, do not necessarily create privacy issues.
2. Although there is much hype currently surrounding AI, data-driven decision making is not new. However, there are new types of deployments of AI. In many areas, there are sufficient laws in place to regulate potential harms, but we are too early with respect to many of these new deployments to fully understand what other potential harms may exist, much less what the legal recourse should be.
3. Because SAS understands that these new deployments of AI involve massively complex models and data sets that potentially contain bias and discrimination, it is in our best interest to take steps to incorporate interpretability, accountability, and fairness into the design and implementation of our models. While many of these models do not necessarily need personally identifiable information to achieve accurate results, all AI systems rely on massive data. The US has sectoral rules relating to privacy and data, but since AI can be deployed to solve the same types of problems in various industries, it may well be time for the US to develop a comprehensive set of rules governing the collection and use of data.

SAS and AI History

SAS has been in the artificial intelligence business for more than 42 years. The use of computerized statistical models to understand data is not new. As a field, AI has been evolving for more than 60 years. In the 1950s through the 70s, we saw advances in the development and use of neural networks; in the 1980s through the mid 2000s, we saw developments in machine learning. Today, with advances in computing power, deep learning, and natural language processing, computers can now have "cognitive abilities"—the ability to understand language and detect objects and pictures. The power of AI has always been that it enhances human effort, but we are seeing that augmentation now in new and exciting forms.

There are many definitions of what AI is, and many subsets or applications of AI. As an initial suggestion, the FTC may want to more granularly define the information that it is seeking by with this particular question by defining key terms. SAS' definition of AI generally is the science of training systems to emulate human tasks through learning and automation. With AI, machines can learn from experience,

adjust to new inputs, and accomplish specific tasks without human intervention. Core AI capabilities for SAS include: machine learning (training the computer to automate learning from complex data sets); natural language processing (teaching computers to understand, interpret and manipulate human language); and computer vision (teaching computers to process, analyze and understand images).

Consumer Welfare Impacts from the Use of Data-Driven Decision Making

Many companies use analytical tools incorporating AI (collectively, analytics) to help them understand not only what their individual customers want, but how to deliver those goods and services to them in the most effective manner. For example, marketing analytics can help companies tailor offers to customers that are most relevant to them. For the consumer, the benefit is receiving an offer of value, while avoiding unwanted and irrelevant offers. Retailers use analytics to help them understand not only which products customers want, but what their channel purchasing preferences might be. Similarly, analytics help them make decisions about the optimal price at which consumers will purchase these goods. Content providers use analytics to make personalized recommendations. With AI, these decisions are increasingly automated, and can surface patterns and trends much more quickly—and from new data sources (such as audio and visual records.) Manufacturers use analytics to detect defects in the manufacturing processes, helping them provide safer products that are desired by consumers. Hotels and other businesses use analytics not only to help them optimally price rooms for customers, but to improve the individual customer experience. Banks and other financial entities use analytics to detect real-time fraud, saving customers from theft and loss, but also to provide identity protection. Medical companies are now using analytics to help diagnose specific medical decisions. While the core of these applications is data, most do not require personally identifiable information to make decisions that benefit individual consumers.

These use cases are not new. The advances in AI, however, improve the model accuracy (discussed further below), can ingest much bigger data, and can operate both more quickly and much more autonomously. It is also not a new phenomenon that sometimes the results are faulty, leading potentially to the denial of goods and services. In many of these areas, regulations currently exist to regulate the collection and use of data, and to prevent abuse and discrimination. For example, the Fair Credit Reporting Act provides recourse if someone is denied access to credit for improper reasons. Similarly, fair housing laws prevent redlining and the improper denial of housing to candidates. The Health Insurance Protection and Portability Act already prevents the improper use and disclosure of sensitive medical information.

The existing sectoral regulations are, for now, sufficient to provide consumer recourse in the face of known harms. In other areas, however, AI is still in early stages of development, and additional regulations, without understanding potential concrete harms, are likely to discourage continuing innovation. For example, there are real concerns that the GDPR restrictions relating to "profiling" will negatively impact innovation for EU consumers. Similarly, the recent passage of a broad privacy law in California including in the definition of "personal information" "inferences" not directly linked to identifiable individuals, is likely to have negative impacts on the ability of companies to continue to innovate with respect to California citizens. While developing specific regulations for some of these AI applications is premature, a bigger risk is the development of competing, disjunctive regulatory approaches. An unharmonized system of regulations likely presents a worse-case scenario for consumer welfare and innovation—some consumers may see the benefits of analytics, while others may not. We see this, for example, as a result of differing approaches to whether data can freely flow across borders. In short, while SAS believes that we are not at a

point where there is benefit to creating additional regulation on analytics and AI, to the extent there may be developing regulations, the FTC should ensure harmonization in this area.

The Developing Analytics Market, Opportunities, Consequences, and Responsibilities of AI Developers

Today, there are two broad categories of "analytics": "white box" techniques, and "black box" techniques. Techniques included within the "white box" category include decision trees, regression analysis, rule lists etc. In general, these techniques, when used with few parameters, are easy to understand. They use interpretable transformations and give users more intuition about how things work, which helps them understand what is going on in the model. Users of these techniques can explain them to a technical audience. But of course, if there are hundreds of features and the modeler builds a very deep, large decision tree, things can still become complicated and uninterpretable.

On the other hand, "black box" techniques include deep neural networks, random forests and gradient boosting machines, which use many predictors and complex transformations. "[T]he complexity that gives extraordinary predictive abilities to black box models also makes them very difficult to understand and trust. The algorithms inside the black box models do not expose their secrets. They don't, in general, provide a clear explanation of why they made a certain prediction. They just give us a probability, and they are opaque and hard to interpret. Sometimes there are thousands (even millions) of model parameters, there's no one-to-one relationship between input features and parameters, and often combinations of multiple models using many parameters affect the prediction. Some of them are also data-hungry. They need enormous amounts of data to achieve high accuracy. It's hard to figure out what they learned from those data sets and which of those data points have more influence on the outcome than the others." *Interpretability is Crucial for Trusting AI and Machine Learning,* Ilknur Kayner Kabul, *Subconscious Musings Blog,* December 18, 2017, available at https://blogs.sas.com/content/subconsciousmusings/2017/12/18/interpretability-crucial-trusting-ai-machine-learning.

SAS is taking several steps to improve the understanding of how these algorithms are functioning. By doing so, we engender both greater trust from our customers, and improve their ability to help their customers understand why particular decisions were or were not made. First, SAS is endeavoring to embed "interpretability" into its analytics and AI. By interpretability, SAS means that it is attempting to give explanations to end users for particular decisions or processes. Stated another way, SAS is trying to help give users "the ability to understand which examples the model gets right or wrong, and how it is coming up with the answers it gets." See *How to Make Your Data and Models Interpretable By Learning From Cognitive Science*, Catherine Olsson, available athttps://medium.com/south-park-commons/how-to-make-your-data-and-models-interpretable-by-learning-from-cognitive-science-a6a29867790.

For SAS, interpretability entails:

- Understanding the main tasks that affect the outcomes.
- Explaining the decisions that are made by an algorithm.
- Finding out the patterns/rules/features that are learned by an algorithm.
- Being critical about the results.
- Exploring the unknown unknowns for that algorithm.

Interpretability also applies to different phases of the modeling exercise. For example, data scientists want to be able to interpret "model inputs". That is, it is critical to understand the data set that one has before starting to model. The old adage "garbage in garbage out" applies with equal force in the AI context. Exploration, cleansing and preparation of data for modeling are critical steps to building accurate models, checking assumptions, and understanding the data that is proposed to be used. A variety of exploratory data analysis and visualization techniques exist to help data scientists understand their data sets before modeling, and most good data scientists will spend the time upfront to understand their data before starting to model.

Second, SAS is incorporating within its products frameworks that will allow our users to evaluate and visually inspect the models to arrive at their own interpretations of how the models and analytics are working. For example, within our latest platform offerings, SAS has begun to incorporate techniques such as Partial dependence (PD) plots, Individual Conditional Expectation (ICE) plots, and Local Interpretable Model-Agnostic Explanation (LIME) protocols. These capabilities enable our customers to visualize complex data, offering a transparent view into how our models react, giving them insights into why these complicated models reacted to certain variables. See *Interpreting Black Box Machine Learning Models Using Partial Dependence Variables and Individual Conditional Expectation Plots*, Ray Wright, Paper SAS 1950-2018, available here https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2018/1950-2018.pdf . PD variables and ICE plots are visual, model-agnostic techniques that depict the functional relationships between one or more input variables and the predictions of a black-box model. For example, a PD plot can show whether estimated car price increases linearly with horsepower or whether the relationship is another type, such as a step function, curvilinear, and so on. ICE plots enable data scientists to drill much deeper to explore individual differences and identify subgroups and interactions between model inputs. *Id.* These techniques do not, by themselves, reveal the inner workings of the model, but instead show how the model behaves in response to changing variables. *Id.* With embedded machine learning attributes, decisions generated by SAS analytics and AI are explainable and accountable.

Finally, in developing SAS products, including those incorporating AI, SAS adheres to several "by design" concepts: fairness by design, accountability by design and transparency by design. *Building Trust in Machine Learning and AI*, Oliver Schabenberger, InfoWorld, February 15, 2018, available at https://www.infoworld.com/article/3255948/machine-learning/building-trust-in-machine-learning-and-ai.html. By adhering to these design principles, SAS is acknowledging that it needs to incorporate ethical principles while addressing societal concerns in developing its analytical solutions. And by incorporating such principles into our design mentality, we have a greater ability to understand where and how bias might enter into the algorithm, and as a result, to "create algorithms carefully designed to avoid or reduce bias—algorithms that can lead to widespread elimination of discrimination." *Id.*

While SAS is taking steps to improve the interpretability of our models, a major issue with respect to transparency for any system relates to data governance. It is axiomatic that "advanced analytics, and AI in particular, need data. What's more, it needs lots of data, and it needs to be relevant and of good quality." *The Bigger AI Challenge—Governing Data,"* Josefin Rosen, *Hidden Insights Blog,* August 15, 2018, available at https://blogs.sas.com/content/hiddeninsights/2018/08/15/bigger-ai-challenge-governing-data. "More data means better trained algorithms, which means more precise and therefore more accurate predictions. This, in turn, means better, evidence-based business decisions." *Id.* But because the models learn everything from the data that is ingested, "[i]f the data is reflecting bias, discrimination or similar, then so will the AI system." *Id.* Data governance is a means of understanding what is in the data and managing the model evolution (and tracking changes in the modeling). One way to view data governance is as a three

step process: define the data that is available and how it can be used; assess what data might be missing; and collect the missing data. *See id.* As part of this process, a variety of exploratory techniques can and should be used to help assess what the data set contains. *See, for example, Understanding and Interpreting Your Data Set*, Ilknur Kaynar Kabul, *Subconscious Musings Blog*, March 9, 2018, available at https://blogs.sas.com/content/subconsciousmusings/2018/03/09/understanding-interpreting-data-set.

While SAS believes that there are sufficient laws regulating the negative impacts that may arise from analytical applications, there are inconsistent rules relating to data. In the US, the rules are patchy and sectoral, which may create gaps for analytics and AI that have multi-industry application. Although many companies have positive incentives to deploy effective data governance and deployment policies and processes, recent events suggest that not all will. It may be time, then, to develop a more comprehensive framework relating to the collection and use of data, particularly for those applications that require both personally identifiable information and sensitive personal information. As a prelude to the development of such a framework, the FTC is in a unique position to help entities understand the criticality of data governance principles and to promote the adoption of effective policies as a best practice.

3. Conclusion

There is a saying at SAS—"We do well by doing good". Understanding, and mitigating potential negative consequences of analytics and AI is built into our dna. And by helping to create responsible algorithms in an ethical, explainable, and accountable manner, we will both continue to do well while continuing to do good. We appreciate the opportunity to provide our perspective on this important topic, and commend the Federal Trade Commission for opening this dialogue. We look forward to being a constructive participant.