August 12, 2024

Mr. Moses Kim
Director
Office of Financial Institutions Policy
Department of the Treasury

Re: Uses, Opportunities, and Risks of Artificial Intelligence in the Financial Services Sector

Dear Mr. Kim:

Block, Inc. ("Block") appreciates the opportunity to submit these comments in response to the Department of the Treasury's request for information ("RFI") on the uses, opportunities, and risks presented by developments and applications of artificial intelligence ("AI") within the financial services sector.

Our Financial Services Business

Block, Inc. (NYSE: SQ) (formerly, Square, Inc.) is a global technology company with a focus on financial services. Our business units (Square, Cash App, Afterpay, TIDAL, and TBD) build tools to help more people access financial services. This letter focuses specifically on the financial services products we offer at Square, Cash App, and Afterpay. At Cash App we have over 57 million[1] active monthly users, at Square we service more than 4 million sellers[2], and at Afterpay we have 24 million active consumers.

Square makes commerce and financial services easy and accessible with its integrated ecosystem of commerce solutions. Square offers purpose-built software to run complex restaurant, retail, and professional services operations; versatile e-commerce tools; embedded financial services and banking products; staff management; payroll capabilities; and much more – all of which work together to save sellers time and effort. Millions of sellers across the globe trust Square to power their business and help them thrive in the economy.

Cash App is a financial services platform that aims to make consumers' relationship with money more relatable, instantly available, and universally accessible. Cash App offers ways to spend, send, save, and store money. Sending and receiving money is free and fast, and most payments can be deposited directly to an

[1] Block Investor Presentation Q1 2024
[2] About Square, "More than four million sellers run on Square, the only platform of its kind designed for solo entrepreneurs and the largest international chains."

external bank account in just a few seconds. With Cash App, customers can also buy and sell stocks and Bitcoin instantly, get a paycheck deposited right to the app, create a unique $cashtag to share with anyone to get paid, and use the Cash App Card to spend their money everywhere Visa Debit is accepted. With these features, Cash App is revolutionizing the way people interact with money, making it simpler and more accessible for everyone.

Afterpay is transforming the way we pay by allowing anyone to buy products immediately and pay over time - enabling simple, transparent and responsible spending. It is on a mission to power an economy in which everyone wins. Afterpay is offered by thousands of the world's favorite retailers and used by millions of active global customers. Afterpay is currently available in Australia, Canada, New Zealand, the United States and the United Kingdom, where it is known as Clearpay.

This Submission

Block believes that the responsible adoption and regulation of AI within financial services yields significant benefits for consumers, businesses, and the broader economy. AI and Machine Learning ("ML") can be used to detect and prevent risk and fraud, promote regulatory compliance, personalize product offerings, improve customer service, and support fair lending decisions. Our experience with AI-driven solutions across Cash App, Square, and Afterpay demonstrates that AI can be used to streamline operations and enhance financial inclusion.

As the Treasury Department gains information on the risks and opportunities associated with the use of AI in financial services, we urge the adoption of a nuanced, principles-driven approach to any AI regulatory proposals. We also recommend opposition to imposing broad and potentially stifling regulations across entire industries. By focusing on addressing genuine risks that may exist due to potential regulatory gaps and maintaining a pro-innovation stance, regulators can ensure that the use of AI continues to responsibly advance, provide increased and efficient access to financial services for all communities, contribute to economic growth, and enhance global competitiveness.

Before any new regulations of AI are proposed to address the risk of AI-driven bias, we recommend that regulatory authorities conduct a gap analysis to determine what aspects of bias mitigation are not fully addressed by existing guidance, laws, and regulations (e.g. the Equal Credit Opportunity Act (ECOA), Fair Credit Reporting Act (FCRA), and the Federal Reserve's Supervisory Guidance on Model Risk Management (SR 11-7)), and conduct a formal information gathering process through extensive collaboration with public and private stakeholders in order to develop consistent standards for bias mitigation.

With these values in mind, our submission addresses the following questions in the RFI: (i) Question 1, the definition of AI; (ii) Question 5, the expected benefits of AI; (iii) Question 7, risk management frameworks for AI; (v) Question 13, the use of AI

in addressing and mitigating illicit financial risk; and (vI) Question 18, policies related to the use of AI.

Question 1, The Definition of AI

Is the definition of AI used in this RFI appropriate for financial institutions? Should the definition be broader or narrower, given the uses of AI by financial institutions in different contexts? To the extent possible, please provide specific suggestions on the definitions of AI used in this RFI.

The definition of AI[3] in the Treasury's RFI, taken from the President's Executive Order on AI, is a broad definition that appears to include many subsets of AI, like ML, that have been used widely in financial services and have been regulated by existing financial regulations for decades. We believe that this definition is too broad, and there should be a distinction between different types of AI that are encompassed under the broader definition.

In short - AI systems run a wide gamut. AI systems include, but are not limited to:

- ML, a subset of AI in which a program or system trains a model from input data.[4] The trained model can make useful predictions from new (never-before-seen) data drawn from the same distribution as the one used to train the model. Machine learning includes techniques like linear regression, logistic regression, decision trees, boosted trees, and neural networks.[5]
- Generative AI, a subset of AI that focuses on creating new content that is complex, coherent, and original,[6] such as Chat-GPT for text generation.
- Robotics, which incorporates AI techniques for autonomous control and navigation, such as self-driving cars and industrial robots.
- General AI, which currently is a theoretical concept that refers to AI systems with the ability to understand, learn, and apply knowledge across a wide range of tasks similar to human intelligence.

The different subsets of AI pose different risks, operate differently, and should be considered and regulated with those differences in mind. For instance, generative AI is a relatively new subset of AI that has a unique set of risks and opportunities that should be addressed narrowly and explicitly. On the other hand, the use of ML has been in existence for decades and is already well understood and regulated in the financial services industry. Additionally, the AI use cases ultimately have different risks associated

[3] "Artificial intelligence" or "AI" has the meaning set forth in 15 U.S.C. 9401(3): a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human--based inputs to perceive real and virtual environments; abstract such perceptions into models through analysis in an automated manner; and use model inference to formulate options for information or action.

[4] Source: https://developers.google.com/machine-learning/glossary.

[5] Additional detail on ML algorithms is included in the Appendix.

[6] Source: https://developers.google.com/machine-learning/glossary.

with them as well. For example, the types of risks and regulations associated with AI in lending are very different from the risks and regulations associated with the use of AI in customer support or recommender systems.

We have seen attempts at the state and federal level to broadly regulate the use of "AI in financial services" by utilizing the all encompassing definition of AI, as used in this RFI and the President's Executive Order. We have concerns with this approach, as it does not distinguish between different subsets and use cases of AI. This runs the high risk of imposing duplicative or overly burdensome requirements on an industry that is already well regulated—such as the use of ML models for credit underwriting in financial services (discussed below). As an alternative approach, we recommend that if regulators or policymakers identify certain risks they are trying to address, they should also identify which subset and use case of AI is related to the risk. Any definitions should be narrowly targeted to avoid overregulation or conflicting regulatory requirements. As AI continues to develop and new innovative tools emerge, more narrow definitions will continue to be of growing importance.

Question 5, The Expected Benefits of AI

How has the use of AI provided specific benefits to low-to-moderate income consumers and/or underserved individuals and communities (e.g., communities of color, women, rural, tribal, or disadvantaged communities)?

How has AI been used in financial services to improve fair lending and consumer protection, including substantiating information? To what extent does AI improve the ability of financial institutions to comply with fair lending or other consumer protection laws and regulations? Please be as specific as possible, including details about cost savings, increased customer reach, expanded access to financial services, time horizon of savings, or other benefits after deploying AI.

How can AI provide benefits to underserved communities?

We believe that AI is a tool that can be wielded to materially increase access to financial services, create innovative products, improve productivity for small businesses and hence spur economic growth. There are a number of ways we utilize AI at Block to increase access to financial services for small business owners and underserved communities.

At Square, we believe our north star for integrating generative AI into our software is how we can help our sellers reduce their time spent on administrative tasks (inventory, billing, menu construction, correspondence, etc.) and enable them to spend more time on their core business and customers. Our Square business has been rapidly building out its toolkit of AI-powered products[7] to help sellers save time and find cost savings with offerings that generate compelling content, drive

[7] See https://squareup.com/us/en/ai-for-businesses.

operational efficiency through automation, and streamline customer and staff communications. For example, through our menu generator, restaurants can gain valuable momentum opening or relaunching by populating their full menu in just a few minutes and with just a few inputs.

At Cash App, we utilize AI for personalization in the Cash App Offers[8] product, which enhances user experience and offers discounts for our users. We offer ways for people to save on every day spending through exclusive discounts with Cash App Pay and the Cash App Card. There are many people who do not have credit cards and therefore cannot earn points or benefits for their purchases. One thing that makes us unique in the financial services industry is that we offer personalized discount offerings for the Cash App debit card so that customers can still experience benefits without having to rely on credit cards that have the potential to turn into revolving debt. For example, a customer can use their Cash Card to get 10% off Uber or a number of various coffee shops, restaurants, and places they already shop. These discounts for everyday purchases provide impactful savings for underserved communities.

Another way we have utilized ML to offer benefits to underserved communities is through our lending products. Block has used ML models to offer lending through Square Loans, Cash App Borrow, and Afterpay. Our ML models have provided significant benefits to low-to-moderate income consumers and underserved individuals by enabling more accurate and fair credit assessments, which, in turn, enables Block to increase our ability to lend to these communities. For example, Square Loans leverages ML to extend credit to small businesses by analyzing real-time sales transactional data, which provides a comprehensive picture of a business's risk and creditworthiness. Square's entirely online offering expands credit access to small businesses that may not otherwise have access to traditional financing options because they lack traditional indicia of credit-worthiness. Since its public launch in May 2014, Square Loans has facilitated more than 2.1 million loans and advances, representing more than $15.1 billion in principal amount loaned or advanced. This includes approximately 150,000 loans to small businesses representing more than $1.5 billion of Paycheck Protection Program ("PPP") loans facilitated in 2020 and 2021, excluding canceled loans.[9]

Our continuous development and improvement of these ML models has enabled Block to offer more and larger loans. Through enhanced risk management and ML models, we have increased our maximum loan size. AI has helped us manage the risk associated with these offers, enabling greater financial inclusion.

Through Cash App's Borrow feature, we offer customers an easy, fair, and accessible way to receive additional liquidity. We utilize ML models that leverage financial data directly indicative of creditworthiness—including real-time spending and savings metrics—to offer credit to underserved populations. Our ML models allow us to make loan offers and determinations quickly, compared to other funding

[8] See https://cash.app/spend
[9] 2023 Block Annual Earnings Report

providers that have lengthy application processes. Our services are also offered through an app, which increases our ability to serve a wide range of customers who don't have easy and recurring access to a traditional brick-and-mortar bank location. The fact that we are able to do this is entirely due to the fact that we make significant use of ML to manage risk. This promotes financial inclusion and provides access to credit for those typically overlooked by traditional credit providers because we do not rely on credit scores to make risk determinations - we consider a wide range of internal variables related to how customers interact directly with our platform.

Similarly, Afterpay uses ML to manage risk of loss, similarly to Square Loans and Cash Borrow. The use of ML enables Afterpay to maintain highly accurate predictions of customer default and fraud risk, which in turn produces efficient credit allocation and helps to maximize financial inclusivity.

How has AI been used in financial services to improve fair lending and consumer protection, and comply with fair lending or other consumer protection laws and regulations?

AI improves fair lending and consumer protection because unlike traditional approaches to determining a consumer's creditworthiness, which are inherently biased, AI expands the universe of predictive attributes into non-biased values. AI enables entities to utilize non-traditional factors to predict creditworthiness, ability to pay, potential for default and the like, without using attributes that are known to carry inherent bias. Block uses a consumer's real-time financial data based on their daily activity through our financial services, such as payment processing, peer-to-peer transactions, debit card purchases, banking activity, etc. This approach significantly expands access to credit for underserved populations and businesses, promoting financial inclusion and equity. Additionally, AI ensures better risk management by dynamically adjusting risk exposure and assessing creditworthiness more accurately. In addition to achieving the critical goal of financial inclusivity, other specific benefits include cost savings and increased systemic safety and soundness through reduced defaults, and early detection of financial stress.

Risk management and compliance oversight benefit significantly from AI, as it enhances the ability to detect fraudulent activity and protect legitimate customers. AI-driven models enable real-time monitoring and dynamic risk assessment, which not only safeguard the financial ecosystem but also ensure the responsible expansion of financial access.

By efficiently managing risks, AI helps maintain the integrity of financial services while widening their reach. For example, in Cash App, we use machine learning models designed to detect P2P payment scams in real-time. These models work through algorithms based on customers' reports of potential scams. Upon receiving these reports, we evaluate and investigate the reports as appropriate to address scam related activity, and that data is used to improve the performance of our models. Other data sources that feed these ML models include previous customer support cases, data associated with our scam warnings, and a host of

additional customer and transaction information. We monitor metrics surrounding how our customers interact with our scam payment warnings and the kinds of payments that trigger those warnings to continue optimizing our deployment of the tool. As these models are fed with the data, they become better and better at predicting and identifying potential scams on the platform.[10]

Existing Regulatory Guidelines that Apply to Lending ML Models

Block's use of ML in our lending businesses is not new or unregulated. Many techniques in ML, such as decision trees and regression, have been applied and subject to regulation long before they were even categorized under the umbrella of AI and instead were referred to as "statistical models", or simply "models" within each regulator's guidance on model risk management such as those in the SR 11-7[11], the Office of the Comptroller of the Currency's (OCC) 2011-12 Guidance[12], or the Federal Deposit Insurance Corporation's (FDIC) FIL-22-2017.[13]

Under existing law, financial institutions utilizing AI/ML are held to a high regulatory standard to mitigate bias, discrimination, or unfair lending practices. Financial institutions must also implement robust model risk management practices, including rigorous validation, ongoing monitoring, and governance frameworks, to ensure safety and soundness. Lenders are subject to examination by U.S. national prudential regulators (e.g. FDIC, OCC, FRB, National Credit Union Association (NCUA)), U.S. national consumer protection regulators (e.g. Consumer Financial Protection Bureau (CFPB), Federal Trade Commission (FTC)), international prudential regulators (where operating internationally, e.g. Prudential Regulation Authority of the Bank of England (PRA), European Central Bank (ECB), Australian Prudential Regulation Authority (APRA)), and state regulatory authorities. Many of these regulators have overlapping scopes that include examination, assessment, and supervision of lenders to ensure the lenders comply with lending regulations (including the use of ML in lending) including: ECOA (Regulation B), FCRA(Regulation V) Truth in Lending Act (Regulation Z), Gramm-Leach-Bliley Act (Regulation P), General Data Protection Regulation (where applicable), California Consumer Privacy Act, Bank Secrecy Act (anti-money laundering regulations), Consumer Financial Protection Act, and the FTC Act.

The components built into SR 11-7, OCC 2011-12, FIL-22-2017 and ECOA, although developed prior to the introduction of AI and ML, form a solid foundation for bias mitigation. Block adheres closely to these principles from both a model/credit risk management and ECOA perspective. For example, ECOA's disparate impact standard requires entities to evaluate models and other algorithmic inputs for proxy values, to gather and test data, and if warranted, consider whether a

[10] See Cash App Blog: "How AI-Driven Scam Warnings Help Prevent Fraud on Cash App" (August, 2024).
[11] See Federal Reserve's SR 11-7 Guidance on Model Risk Management.
[12] See OCC's 2011-12 Guidance on Model Risk Management.
[13] See FDIC's FIL-22-2017 Guidance on Model Risk Management.

less discriminatory alternative would satisfy a legitimate business need while achieving the desired business outcome. SR 11-7 incorporates several bias-reducing mechanisms; indeed, all four governing principles of SR 11-7 work together to prevent, detect and mitigate bias: model development, validation, implementation, and use of models. The two primary sources of model risk identified by SR 11-7 are specific to unwanted and socio-economically damaging model outcomes: (1) fundamental errors and production of inaccurate outputs when viewed against the design objective and intended business uses; and (2) incorrect or inappropriate use of models.

Question 7, Oversight of AI – Explainability and Bias

What challenges exist for addressing risks related to AI explainability? What methodologies are being deployed to enhance explainability and protect against potential bias risk?

Potential Challenges in AI Explainability

AI explainability is a nuanced topic that might not be relevant to all applications of AI; in some cases, the nature and scope of the AI's use will dictate whether explainability is essential. Different applications have distinct regulatory requirements and necessitate tailored methodologies for explainability. It is crucial to consider the specific context and application when evaluating the need for explainability, as this can significantly impact how methodologies should be designed and implemented. At Block, we recognize that there are risks related to AI explainability in the financial services sector, particularly when it comes to lending and mitigating bias. Below, we focus on identifying those general risks faced by the industry and the methodologies to deploy to enhance explainability and combat the risk of bias, specifically in lending products.

Complexity and Opacity: Some AI models might be highly complex and sometimes considered "black boxes." This complexity can occasionally make it difficult for stakeholders to understand how input features translate into specific decisions.

- **Regulatory Compliance:** Given regulations like the ECOA and the FCRA, there might be challenges in clearly communicating reasons for adverse actions (e.g., credit denials) from complex models to consumers. Different AI applications may face varied compliance requirements, necessitating careful consideration of which explainability techniques are appropriate and how they should be applied to ensure clear and compliant communication.
- **Potential Bias:** Some AI models might inadvertently incorporate biases present in the training data, potentially leading to discriminatory outcomes for

protected classes. Although rare, ensuring fair treatment while maintaining model performance remains a crucial, albeit manageable, task.

Methodologies to Enhance Explainability and Mitigate Bias

While we implement various methodologies to enhance explainability and mitigate bias across multiple aspects of our AI practices, it's important to recognize that these approaches may not be universally mandated or applicable for all AI applications. Each use case should be carefully assessed to ensure that explainability, fairness, and compliance requirements are appropriately addressed. The following methodologies are commonly utilized in our consumer lending models but may vary across different AI use cases to better meet specific needs and challenges.

- **Fair Lending Compliance Review Prior to Model Deployment.** In the context of consumer lending, all models should adhere to fair lending laws and consumer protection regulations. This review process identifies and addresses potential compliance issues before models are deployed. By rigorously analyzing models for any discriminatory impacts, these practices should align with regulatory expectations and uphold the principles of fairness and transparency. For other AI applications, such as recommendation systems used for personalized marketing, compliance reviews might focus more on data privacy and ethical considerations rather than specific lending laws.
- **Excluding Variables Associated with Protected Class Status.** To mitigate bias risk and comply with regulations in consumer lending, variables directly related to protected class characteristics (e.g., race, national origin, age) are excluded from models. This best practice of prohibiting the use of such variables and setting stringent guidelines on other sensitive factors like location or age is used to prevent unintended biases. This proactive approach can help ensure that models are fair, compliant, and respect individual privacy and rights. In other AI use cases, such as marketing and ad targeting, demographic information might be necessary for delivering relevant content, thus the focus would be on ethical usage and clear communication rather than exclusion.
- **Applying ML Explainability Techniques.** In consumer lending, various ML explainability techniques can be employed to interpret model outputs and facilitate understanding for both businesses and customers. Techniques such as SHAP (SHapley Additive exPlanations), LIME (Local Interpretable Model-agnostic Explanations), and Partial Dependence Plots (PDP) help break down model decisions into understandable components. SHAP values attribute a score to each feature's contribution, LIME approximates complex models with simpler, interpretable models, and PDP provides insights into

how specific features impact predictions. Leveraging these techniques ensures that adverse action notices (AANs) are clear, making credit decisions transparent and comprehensible. For other AI applications, such as fraud detection systems, explainability might involve providing sufficient information to understand why a transaction was flagged as suspicious, focusing on reducing fraud rather than detailed transparency to the end-user.

Reliance on SR 11-7 for Model Development and Model Risk Management

Model development and risk management strategies in consumer lending align with SR 11-7 guidelines, which should be a cornerstone for any institution in the financial sector. These best practices are reflective of the stringent regulatory environment in consumer lending due to the risk of disparate impact and associated compliance requirements. For other applications of AI, tailored strategies may be more appropriate to meet their specific needs and regulatory contexts.

- **Robust Model Documentation.** In consumer lending, detailed documentation of model design, data inputs, and assumptions should be maintained to enhance interpretability and traceability. This practice aligns with SR 11-7 and ensures models are understandable and transparent. For other AI applications, comprehensive documentation should be adapted to reflect the unique aspects of each use case to support transparency and interpretability.
- **Validation.** Validations in consumer lending should be conducted to confirm that models function as intended, with conceptual soundness reviews and performance testing. This approach follows SR 11-7 guidelines to validate model efficacy and reliability. For other AI applications, validation processes should be customized to address the specific risks and requirements of those models.
- **Ongoing Monitoring.** Continuous monitoring of models post-deployment in consumer lending should occur, which allows for the prompt identification and addressing of emerging risks. This ongoing vigilance ensures alignment with SR 11-7's requirements and maintains model integrity over time. For other AI applications, ongoing monitoring should be adjusted to match the particular risk profiles and operational dynamics of those models.
- **Risk Based Approach.** The extent of investment in each of these practices is a function of the model's degree of risk and the relevant regulations and concerns. Higher-risk models and those subject to stricter regulatory environments, such as those in consumer lending, demand more rigorous documentation, validation, communication, and monitoring. In other contexts with different risk levels and regulatory requirements, investments should be

proportionately tailored to balance effective model governance with operational efficiency.

Question 13, Using AI to Detect and Mitigate Fraud

How do financial institutions, technology companies, or third-party service providers expect to use AI to address and mitigate illicit finance risks?

AI is a crucial component of mitigating risk and preventing fraud in our business segments.

Square

Square Risk Manager is a customer-facing software application that we developed to identify unusual payments and fraud patterns for our Sellers.[14] It leverages ML to spot and block online payment fraud and provides real-time transaction monitoring to protect businesses of all sizes.

Cash App

Our Cash App business has integrated AI models for risk monitoring to ensure that accounts are accessed only by valid account holders, to prevent the use of stolen instruments and identities, and to block fraudulent transactions. Our team identifies key signals to collect, metrics to track, and actions to take.

Cash App maintains risk ecosystems across several event types. Each ecosystem consists of many (sometimes as many as several-hundred) ML models. The purpose of each model is to review specific events, and to assign a value to that activity which approximates the likelihood that the event is associated with unwanted activities. If the resulting value for any model exceeds its predetermined threshold, we will flag the event for action. The flagging thresholds for each model are systematically determined to provide the best performance for the given ecosystem. In order to minimize false positives while also reducing unwanted activity, we target specific problem spaces. The problem spaces we target include transactions resulting in financial loss through chargebacks and disputes, account takeovers, scam transactions, banking fraud, and stolen identities.

The AI system operates by regularly training models with a wide set of signals to capture the specific fraudulent behaviors of the current bad actors. We train new models very frequently, and adjust the thresholds of our models to ensure that we're capturing the current behavior with the highest precision and recall. The largest issues we run into are over-governance of our models, resulting in slow turnover of our models, resulting in an inability to protect our customer base.

[14] See https://squareup.com/us/en/payments/risk-manager.

Afterpay

For our Afterpay business, we have developed models to prevent various types of fraud and abuse. These models predict when transactions might come from unauthorized users, when customers might be using stolen financial instruments, and when new users might not pay. Each model helps us decline fraudulent transactions and protect our customers.

What challenges do organizations face in adopting AI to counter illicit finance risks?

Organizations face challenges in adopting AI, such as ensuring data quality, regulatory compliance, and overcoming technical expertise gaps. These challenges are not unique to illicit finance but are part of establishing a robust framework for the use of AI. AI introduces new risks, including potential false positives/negatives, biases, and security vulnerabilities. However, with robust data management, continuous staff training, stringent regulatory alignment, and maintaining a balance between AI and human oversight, we are well-equipped to manage these risks and harness AI's full potential in various applications, including combating illicit finance.

Question 18, Policies and Regulations Related to the Use of AI

What actions do you recommend Treasury take, and what actions do you recommend others take?

What, if any, further actions are needed to protect impacted entities, including consumers, from potential risks and harms?

Please provide specific feedback on legislative, regulatory, or supervisory enhancements related to the use of AI.

Block believes that if Treasury and other authorities take any actions relating to AI, they should: (1) recognize that AI is a valuable tool that can be used to benefit both the industry and consumers; and (2) target specific use cases where regulatory gaps may exist.

As we alluded to under Question 1, we have seen legislation arise at the state and federal level to broadly regulate the use of AI in financial services that would overregulate an already well-regulated industry. We also have seen attempts to mitigate bias, by requiring industries to revert back to human review, which is inherently biased. For these reasons, we recommend that any new legislative, regulatory, or supervisory approaches take a very nuanced and targeted approach to identifying and mitigating certain risks associated with AI. This requires first defining AI accurately, and then distinguishing between different subsets of AI. Various subsets of AI have varying levels of risk, as do certain use cases.

It is important to conduct a gap analysis and first identify where regulatory gaps may exist when applied to emerging AI technology. It should be noted how existing AI-powered tools, such as the use of ML in credit underwriting, have been used and regulated for decades. Any gaps identified when it comes to fair lending and bias mitigation should be addressed specific to the type of AI being used and the use case - new regulation should not be applied with a broad brush across entire industries, or evenly across all subsets and use cases of AI.

Block also believes that the Treasury Department and its counterparts should leverage the detailed analysis and recommendations NIST has already developed, as it focuses specifically on emerging technologies, rather than introducing an entirely new framework. Block emphasizes that any enhancements made to existing governance principles be designed to enable responsible agility and fluidity within the industry in the face of a quickly evolving technological landscape. Enhancements that are unduly burdensome or redundant will slow the industry's ability to offer new products and services, causing the negative consumer impact and potentially limiting access to AI-powered benefits such as credit risk and offer determinations. Although there are certainly risks that can be identified and mitigated as it relates to AI, there are also many benefits and opportunities that should be allowed to continue to develop in a susceptible regulatory environment - especially when it comes to significantly improving and expanding access to financial services.

We urge the U.S. government to maintain a regulatory environment that supports innovation and global competitiveness, while taking a targeted, principles-based approach to regulating AI. This approach will help ensure AI's responsible growth, all while expanding access to financial services for underserved communities, mitigating risk and illicit activity, and boosting economic growth.

The following are examples of policies that should be avoided as they would impair the development of responsible AI:

- An overly broad definition of AI that does not take into consideration varying risks and opportunities of different subsets of AI.

- An overly broad definition of regulated processing activities (e.g., "automated decision tools") that inadvertently encompasses legitimate ML activities beneficial to consumers, such as fraud prevention and the security of transactions.

- Regulatory frameworks at the state or federal level that require regulatory approval, licensing of AI, and/or assessments prior to the use of AI.

- New regulations that are duplicative of, or conflict with existing laws and regulations, such as FCRA, UDAAP, and ECOA.

- High regulatory compliance costs that are overly burdensome due to the frequency of assessments and reporting requirements.

- Inflexible or disproportionate requirements that may stifle legitimate innovation and reduce competition while requiring entities to comply with regimes that are not suited to their specific activities.

- Providing consumers with the right to opt-out of the use of AI/ML in credit offers or creating a private right of action.

Summary and Conclusion

In this submission, Block aims to highlight how AI is the technology that underpins the financial services industry's ability to prevent, detect and mitigate fraud and illicit activity, and expand credit opportunities. We believe that AI provides a set of critical tools in financial services to the benefit of customers and underserved communities. In fact, there are many existing AI use cases that are now industry standard and well regulated – including, for example, in credit underwriting. But existing and potential uses of AI are not the same and neither are their risks. For that reason, we believe AI should continue to be regulated based on particular uses and the resulting risks.

New regulations should not apply to entire industries or to broad categories of AI technologies, but should be tailored to address use cases where regulatory gaps may exist. In other words, AI policies and regulations should take a tailored principles-driven approach that mitigate risks posed by particular uses while continuing to permit the innovation that allows providers like Square, Cash App, and Afterpay to continue to expand the availability of critical financial services.

Block greatly appreciates the opportunity to provide our insights and we look forward to continued collaboration with the Treasury Department to help shape a regulatory environment in the United States that promotes innovation, fairness, and security in the financial services sector. As further information gathering is conducted related to this topic, please feel free to contact me at the information below or Janessa Lopez, Block's AI Policy Lead, at janessalopez@squareup.com.

Respectfully submitted,

Amanda Anderson

Amanda Anderson

Global Head of Public Policy

amandaanderson@block.xyz

Appendix

Machine Learning Algorithms[15]

The majority of machine learning applications fall into two categories: Regression and Classification. Regression involves predicting a continuous outcome based on input variables, such as forecasting house prices or predicting temperature. Classification involves predicting a categorical outcome, assigning inputs into predefined classes or categories, such as identifying whether an email is spam or categorizing an image as either a cat or a dog.

Linear regression is a statistical method used to predict the value of one variable based on the value of other variables. It does this by finding the best-fitting straight line (or linear relationship) through a set of data points. Essentially, it helps in understanding how the dependent variable (the one you want to predict) changes in response to the independent variables.

Logistic regression is a model used for binary classification problems. It uses a logistic function to model a binary dependent variable, meaning it predicts one of two possible outcomes. It is a type of supervised learning algorithm, where the model is trained on labeled data (data paired with the correct output).

Linear and logistic regression, although predating the term "machine learning," are fundamentally machine learning algorithms. Both models take input data and learn functions that map these inputs to outputs—linear regression predicts continuous outcomes with a best-fit line, while logistic regression predicts categorical outcomes with a separating curve. These models involve training on historical data to optimize parameters, using methods like gradient descent, similar to other machine learning techniques. They also generalize from training data to predict on new, unseen data, a key feature of machine learning. Despite the evolving terminology, linear and logistic regression adhere to the core principles of machine learning: learning from data, mapping inputs to outputs, and optimizing for accurate predictions.

A decision tree is a flowchart-like structure used for classification and regression tasks. Each internal node represents a "test" on an attribute, each branch represents the outcome of the test, and each leaf node represents a class label or decision. It is a supervised learning technique.

Ensemble models are machine learning techniques that combine the predictions of multiple smaller models to improve accuracy and robustness. By aggregating the outputs of these "weak" models, the ensemble model leverages their collective strength, often outperforming any single model. Boosting is a specific type of ensemble method that builds models sequentially, with each new model correcting errors made by the previous ones. This results in a highly accurate predictive model.

---

[15] For more information, see IBM's site on What is a Machine Learning Algorithm?

Gradient Boosting Machine (GBM) is a type of boosting algorithm, which is an ensemble method used for regression and classification problems. It builds an ensemble of weak prediction models, typically decision trees, sequentially, with each new model correcting the errors of the previous ones.

GBM Regression uses Gradient Boosting Machine (GBM) specifically for regression problems. It predicts a continuous outcome using decision trees, correcting the errors of the previous models in a sequential manner.

Light GBM (LGBM) is a specific implementation of the gradient boosting framework, designed to be highly efficient and fast. It uses tree-based learning algorithms and is optimized for both speed and memory usage. Key features include support for parallel and graphics processing unit (GPU) learning, which allows it to handle large-scale data more effectively than traditional gradient boosting algorithms.

XGBoost (Extreme Gradient Boosting) is another specific implementation of the gradient boosting framework, optimized for efficiency and scalability. It enhances the basic gradient boosting algorithm by introducing improvements such as regularization, which helps prevent overfitting, and by employing advanced optimization techniques. As a result, XGBoost excels in terms of predictive accuracy and computational performance, making it a popular choice for machine learning practitioners.

Adaptive learning in machine learning refers to systems that automatically adjust their model parameters based on new data and environments, improving their performance over time without needing manual retraining. Similarly, reinforcement learning is a type of adaptive learning where an agent learns to make better decisions by interacting with an environment and receiving feedback in the form of rewards or penalties. Both concepts highlight self-improvement: adaptive learning models adjust to new data, and reinforcement learning agents fine-tune their actions based on feedback, eliminating the need for manual retraining in these frameworks.