Upstart

Written Statement of
Dave Girouard
CEO and Co-Founder, Upstart Network, Inc.

U.S. Senate AI Insight Forum: High Impact AI

November 1, 2023

Leader Schumer, Senators Rounds, Heinrich and Young, thank you for the opportunity to participate in today's AI Insight Forum on High Impact Artificial Intelligence (AI).

My name is Dave Girouard, co-founder and CEO of Upstart, co-headquartered in San Mateo, California, and Columbus, Ohio with offices also in Austin, Texas. I founded Upstart almost 12 years ago in order to improve access to affordable credit through the application of modern technology and data science.

Upstart is the leading AI lending marketplace, connecting millions of consumers to 100 banks and credit unions that leverage Upstart's AI to deliver superior credit products. With Upstart AI, lenders can approve more borrowers at lower rates across races, ages, and genders, while delivering the exceptional digital-first experience customers demand. More than 80% of borrowers are approved instantly, with zero documentation to upload.

To say recent advances in AI have struck a chord would be an understatement.

For decades, AI languished in academic circles and research labs. We saw glimpses of it in consumer-grade speech-to-text technology developed in the 1990's. As the Internet flourished, we suspected there was something special powering those uncanny Google search results or those spot-on Netflix recommendations. We raised an eyebrow for a brief moment when we learned that the best chess player in the world was no longer a human. Every once in a while, AI found its way to Hollywood screens in the form of dystopian tales sprung from the minds of creative giants such as Stanley Kubrick and James Cameron.

But it wasn't until last year's launch of OpenAI's ChatGPT that public awareness of AI finally escaped the bonds of both academia and Hollywood and entered the mainstream. Suddenly we weren't just talking about board games and fantasy worlds. Anyone paying attention quickly realized something had been unleashed that had the potential to change our world immeasurably and irreversibly.

Generative AI and the large language models that power them woke us up, but AI has been in the works for a very long time. There is almost universal agreement that the potential for AI to do good for humankind is vast. Almost unimaginably vast. But there's also consensus that some bad things could happen. Including some very very bad things.

As we contemplate how best to tap all that goodness that AI can offer, while avoiding all the downside, we're left yearning for a framework that will help us carve a path forward. How much can be left to the private sector and the markets that power it? What type of regulation might be necessary and who's in the best position to lay down the rules? How can America continue to maintain its global leadership position in technology? And what good will any of this do if the terminators of our nightmares are unleashed by bad actors around the world?

And then of course there's this issue of time. Just a few years ago, a smart guy like Nick Bostrom could write a thoughtful book about AI and its dangers. But now we hardly have time to even read a 352 page book, much less write one. The rapid advancements in AI have taken the "do nothing" option off the table. From a regulatory perspective, the Great Financial Crisis and our government's timely and urgent response to it may be the best proxy to the level of urgency required today.

But AI is not a single notion. It's a very broad-based idea that computers can learn in the same way humans learn – by interacting with the real world. AI implies machines can learn by observing and understanding patterns in data that repeat themselves, which is a reasonable way to describe the human mind itself. But in our case, the data is our memories.

Because you're here today, I know you've already thought a lot about the potential – both good and bad – for AI. But I didn't come here to scare you – quite the opposite in fact. I come here to share a compelling example of how AI has been successfully deployed – safely, at scale, in a heavily regulated industry, and with exceptional results. All through a cooperative multi-year effort that involved both the public and private sectors. It is the application of AI to lending, and it is already enabling a credit system that is dramatically more affordable and inclusive.

I know this because Upstart, the company I founded almost 12 years ago, pioneered the application of AI in lending. By way of background, I spent more than eight years at Google, building the foundations of what became the company's $30 billion cloud business. There was a time when cloud computing was almost as scary as AI is today – so I know something about living on the edgier frontier of an emerging technology.

I left Google in 2012 because I wanted to apply the types of technologies we had developed at Google to a different domain – lending and access to credit. My co-founders and I had an

intuitive sense that better math and data science could create more accurate financial risk models. And that those risk models, in turn, could reduce the price of borrowing for many millions of Americans. Some simple numbers illustrated the problem: only about half of Americans have credit scores that would qualify for bank-quality credit. Yet more than 80% of Americans have never actually defaulted on a loan. That sounded like bad math.

So we were focused on "better math," which is how we thought of it in the early days. It wasn't until years later, when the risk models reached an appropriate level of sophistication, that we began to use the terms *machine learning* or *artificial intelligence*. Today, some of the most sophisticated AI model forms in our system – neural networks – are close cousins of those powering generative AI.

While we weren't experts in banking regulation at the time, we at least knew what we didn't know. Our General Counsel was the fifth person to join the company, and she schooled us quickly in the regulations surrounding lending – from consumer protection including ECOA and FCRA, to safety and soundness, to anti-money laundering. We were proposing to transform something right at the heart of banking with a technology that not long ago lived almost exclusively in research labs and we wanted to do it right.

We immediately faced many of the same questions and issues that bring us together today: Will AI-powered lending be fair? Will it introduce bias? Are its outputs and decisions explainable or is it a black box? Can it be managed responsibly from a lender's perspective?

In 2012, before even launching the company, we naively marched up to the San Francisco office of the Consumer Financial Protection Bureau (CFPB) and introduced ourselves. This was not one of the new "Offices of Innovation" – this was the local enforcement team. But what did we know? We were convinced that we were the good guys and were committed to innovating within the law.

In the subsequent years, we worked weekly – and even daily – with the CFPB to determine how AI models could be responsibly applied to lending from a consumer perspective. We developed and refined rigorous models for testing every loan and in fact every single applicant for bias. We shared, and continue to share, the results of these tests with our lending partners as well as with regulators as needed or requested. We developed state-of-the-art methodologies for explaining the outcomes of our models – which in the world of lending are called adverse action notices. We signed a "no action" letter with the CFPB which we maintained across three different administrations related to the application of AI to lending. We also worked as closely as possible with the prudential bank regulators to ensure our platform supported the foundations of lender safety and soundness, including model risk management and third-party risk management. We

made sure lenders on the Upstart platform had complete control over and oversight into the loans they were originating.

Ten years later, America is the world leader in the application of AI to lending. For our part, we work with more than 100 banks and credit unions across the country, who together have originated more than $34 billion in AI-powered loans. We're applying AI to consumer installment loans, small-dollar relief loans, auto purchase and refinance loans, and home equity loans. AI today has allowed us to serve over 2.7 million Americans and has enabled more than 87% of our loans to be approved in an instant – without documentation, phone calls, or waiting. More than 70% of loan applications come from a mobile phone.

Today I leave you with the most important lessons we've learned on our journey at Upstart to use AI responsibly and help establish America as the global leader in AI-enabled lending:

- **AI works** – Because of Upstart's AI, our bank partners can approve at least 40% more borrowers than traditional underwriting methods for bank-quality loans. More specifically, the AI-enabled model approves 43% more Black borrowers than a credit score only model, at 24% lower APRs. Results are similar for Hispanic borrowers.

- **Bias can be avoided** – By rigorously testing every applicant and every loan, we can avoid introducing bias into the system. New versions of our AI models are tested for bias before they are ever put into production. The data from this rigorous testing is validated by third parties, and is shared with each lending partner and with regulatory agencies as needed or requested. Statistically rigorous and standardized testing along with transparent sharing of results can ensure that AI is a win for all.

- **AI can be explainable and accountable** – AI doesn't have to be a black box. At Upstart, we have developed and deployed techniques (referred to as SHAP Values) that clarify the most important contributors to the model's conclusions. These explanations, which are provided to every declined applicant, are both simple to understand and actionable. Similar techniques can be utilized for any decision-making AI application.

- **Clear regulation and jurisdiction helps** – We benefited from relative clarity in consumer protection laws as well as bank safety and soundness guidance. In the case of AI in lending, no new legislation has been passed, but regulators have made efforts to provide useful guidance on interpretation of existing rules. Assuming the use of AI in other domains will require legislation, market participants will benefit from simple and universal rules of the road and clear jurisdiction.

- **Private and public sectors can work together constructively** – Progress for AI in lending was the result of combined efforts from several public, private, and not-for-profit entities. We are members of the Office of the Comptroller of the Currency's Project REACh, the National Community Reinvestment Coalition's Council for Financial Inclusion, and a founding member of More Than Fair – a community of organizations

dedicated to improving access to affordable and inclusive credit for American consumers and small businesses.

In many ways, Upstart benefited from the fact that lending was already so heavily regulated, with ECOA having been law since the 1970s. The principles of safe and responsible lending were already clear and fortunately, they applied reasonably well, even to an AI platform developed almost half a century later.

Because we operate in the narrow domain of lending, and in an area where computers were already more capable than humans, there are other issues associated with AI today that we were fortunately absolved from considering: intellectual property, job and employment dislocation, misinformation, and even the existential risks that are the subject of much debate.

In the broader scope of AI, there is undoubtedly new regulation required to capture the benefits of this rapidly emerging technology while mitigating the risks. But the lessons learned from our decade of leveraging AI to improve access to credit for millions of American families are a good place to start.