COLUMBIA | SIPA
Sustainable Investing Research Initiative

SUMMARY

Columbia University Sustainable Finance Forum

SFF1:
Two Catastrophes, One Blind Spot

September 10, 2026

OVERVIEW

The September 10 Sustainable Finance Forum examined a rapidly emerging question for sustainable finance: **how should investors respond when risks such as artificial intelligence (AI) and climate change threaten not simply individual companies, but entire economies and investment portfolios?**

The session brought together three different perspectives. **David Williams-King**, an AI safety researcher, described the potentially catastrophic risks associated with increasingly capable artificial intelligence. **Christopher Wright**, Head of Sustainability and Geopolitical Risk at Norges Bank Investment Management, described how one of the world's largest diversified investors approaches systemic risks through ownership, corporate engagement, and portfolio stress testing. **Paul Rissman**, co-founder of Rights CoLab, argued that the legal framework governing retirement funds should be changed so that fiduciaries are explicitly required to consider and mitigate systemic risks.

The conversation that followed focused particularly on the limits of investor action, the incentives facing institutional investors, the role of government and law, the difficulty of measuring systemic risk, and whether large, diversified asset owners could become important agents for managing risks that individual companies have neither the capacity nor the incentive to solve.

AI SAFETY: A RISK UNLIKE TRADITIONAL FINANCIAL RISKS

David Williams-King opened by arguing that advances in AI create a category of risk that is difficult to compare with conventional financial or technological risks. AI systems are becoming increasingly capable, can transmit information and learned capabilities much more rapidly than humans, and are increasingly able to act autonomously in the world. At the same time, researchers still have a limited ability to understand how advanced models arrive at decisions or to guarantee that their objectives will remain aligned with human intentions.

Williams distinguished between several kinds of AI risk. One is **misuse**, in which human beings deliberately use AI for harmful purposes such as cyberattacks, disinformation, financial crime, or biological weapons. More difficult is **misalignment**, in which an AI system itself develops or pursues instrumental objectives that were not intended by its designers. An otherwise apparently sensible objective can produce dangerous behavior if a sufficiently capable system discovers that acquiring resources, avoiding shutdown, circumventing restrictions, or taking other unintended actions helps it achieve that objective.

His central warning was that the absence of previous catastrophic AI failures should not be treated as evidence that such failures are impossible. Unlike familiar risks, humanity has no historical experience with systems more capable than ourselves across many cognitive domains. Waiting for conclusive empirical evidence may therefore mean waiting until after a catastrophic event has occurred.

Williams-King also emphasized the **collective-action problem** surrounding AI development. Individual companies may recognize serious safety concerns while simultaneously believing that slowing development would allow competitors—whether other companies or other countries—to move ahead. Competitive incentives can therefore encourage firms to continue developing increasingly powerful systems even when their own researchers or executives believe those systems may present substantial risks.

The implications for investors were left initially as an open question: if the companies creating advanced AI are caught in a competitive race that makes unilateral restraint difficult, could investors, regulators, or other institutions alter those incentives?

THE UNIVERSAL OWNER PERSPECTIVE

Christopher Wright approached AI and climate change from the standpoint of Norges Bank Investment Management, manager of Norway's sovereign wealth fund. Because the fund owns

shares in thousands of companies across countries and sectors, it cannot escape risks that damage the global economy as a whole. Such investors are often described as **universal owners**: diversification protects them against many company-specific risks, but not against systemic risks that affect markets broadly.

Climate change provides an increasingly visible example. Wright noted that physical climate risks are no longer primarily theoretical future possibilities. Droughts, floods, fires, heat, and other extreme events are increasingly affecting economic activity. These physical consequences reflect accumulated emissions from earlier periods and therefore cannot simply be reversed through short-term portfolio decisions.

AI presents a different but also potentially systemic challenge. Its significance is already visible in the concentration of market capitalization among technology companies, in the transformation of financial analysis and investment management, and in the possibility that inadequate safety controls could produce much broader harms. At the same time, Wright stressed that AI also has potentially enormous benefits, including applications in medicine and scientific research.

For Norges Bank Investment Management, one important response is *active ownership*. Large shareholders can vote at annual meetings, engage boards and senior management, establish public expectations, and ask companies to explain their governance and risk-management practices. With AI, this includes asking whether boards are accountable for AI risks, whether companies have adequate safety expertise and resources, whether their systems are sufficiently transparent and explainable, and whether they engage constructively with regulators and standard setters.

A second tool is **stress testing**. Rather than attempting to predict the most likely future, Norges examines low-probability, high-impact scenarios and asks how severely the portfolio could be affected. These scenarios are translated into changes in financial variables and then used to revalue the portfolio. Recent examples included an AI-related market correction, geopolitical fragmentation, sovereign-debt stress, and simultaneous crop failures.

The purpose of such exercises is not necessarily to produce a precise forecast. They make large and sometimes unfamiliar risks visible, quantify the potential magnitude of losses, and allow investors and policymakers to ask whether existing strategies remain appropriate.

Paul Rissman then shifted attention from what investors **can** do to what the law tells fiduciaries they **must** do.

His starting point was that retirement funds are unusually well situated to care about systemic risks. They are large, diversified, and long-term investors. Because they effectively own substantial portions of the economy, damage to the economy ultimately damages their beneficiaries. Yet traditional fiduciary law has concentrated primarily on protecting beneficiaries from **idiosyncratic risk**—risk associated with individual securities—rather than risks affecting the entire market or economic system.

Rissman argued that this produces a major gap. Fiduciaries may in principle be permitted to consider systemic risks, but the legal guidance, measurement methods, and institutional practices necessary to do so remain inadequate. Trustees have well-developed methods for balancing conventional investment risk and return but lack comparable tools for determining how an investment or stewardship action changes portfolio-wide exposure to climate change, AI risk, or other systemic threats.

He connected this problem to the evolution of fiduciary law itself. Earlier reforms incorporated modern portfolio theory into trust law, encouraging diversification and risk-adjusted return rather than minimizing the risk of each asset independently, which led to the unnecessary concentration of trust portfolios in the least risky and lowest-returning assets. Rissman suggested that another intellectual and legal update may now be necessary: from a framework centered predominantly on individual-company and market volatility toward one that also recognizes the probability and magnitude of systemic loss.

His proposed approach would explicitly extend fiduciaries' responsibility for avoiding large losses to systemic as well as idiosyncratic risks and introduce a precautionary principle for circumstances in which severe risks cannot be estimated with conventional precision. The Retirement Protection Task Force is developing model legislation along these lines, including requirements that trustees investigate material portfolio-level systemic risks, consider reasonable mitigation measures, educate trustees and beneficiaries, and undertake systemic-risk stress testing.

Rissman also stressed the need for better analytical tools. Among the possibilities discussed were system-level scenario analysis, attribution of systemic harms to individual companies, portfolio-impact analysis, expected-shortfall measures, and consideration of how risks affect different generations of beneficiaries.

The larger implication of his argument was potentially significant for corporate governance. If retirement trustees were legally required to protect diversified portfolios against systemic risks, their expectations of corporate managers could change. Corporate governance would then need to take account not merely of the returns produced by an individual company but also of the effects that the company's behavior imposes on the rest of an investor's portfolio.

DISCUSSION: CLIMATE AND AI ARE BOTH SYSTEMIC—BUT DIFFERENTLY SO

The discussion brought out an important difference between the two principal risks examined during the session.

One participant, speaking from experience as a public-pension trustee, observed that **climate change and AI create different incentive structures for investors**. Climate damage largely represents a portfolio-wide cost. AI, by contrast, can simultaneously create systemic dangers and extraordinary returns for investors who own the companies benefiting from the technology. Institutional investors may therefore face a particularly difficult conflict: the investments contributing to potential systemic risk may also be among their strongest-performing assets.

That observation sharpened a theme running through the entire conversation: **systemic-risk management is not simply a problem of recognizing risk. It is a problem of incentives**.

The discussion also highlighted trustee capability. Even if the law gives retirement trustees authority to consider climate change or AI, many trustees may lack the technical background necessary to evaluate complex financial, technological, and scientific questions. Education therefore becomes an essential complement to changes in fiduciary law.

Participants also asked what investors can realistically accomplish without complementary government action. Wright emphasized that even exceptionally large investors remain only one part of a much larger system. Investor action can influence companies, but major collective-action problems generally require alignment among investors, governments, regulators, companies, and consumers. Experience with climate policy illustrates the problem: investors may position themselves in anticipation of regulatory change, but if governments fail to implement the expected policies, those investment assumptions may not hold.

Rissman responded that the enormous capital requirements of AI development may give institutional investors unusual leverage. Pension funds and other large asset owners supply both equity and debt capital to technology companies. In principle, they could condition that capital

on stronger AI-safety commitments—for example, by engaging equity holdings, voting against directors, or declining to purchase debt unless companies met agreed safety standards.

AN EMERGING QUESTION FOR SUSTAINABLE FINANCE

Taken together, the presentations suggested a widening conception of sustainable finance.

Traditional investment analysis asks whether environmental, social, technological, or governance factors create material **risks for a company**. This conversation repeatedly raised a different question: **what happens when companies themselves contribute to risks that threaten the economic system on which a diversified portfolio depends?**

Climate change already makes this problem visible. AI may present an even more compressed version of it because technological development is moving rapidly while the relevant governance structures, safety practices, legal doctrines, and measurement tools remain unsettled.

The September 10 Forum did not resolve whether investors can or should become primary agents for addressing these risks. Instead, it identified several interdependent requirements: better understanding of emerging systemic risks; improved methods for measuring their portfolio consequences; better-informed trustees and boards; fiduciary frameworks capable of recognizing portfolio-level harms; stronger corporate governance and stewardship tools; and public regulation capable of addressing collective-action problems that markets alone may not solve.

Perhaps the clearest point of convergence was that **diversification does not protect an investor from everything**. When risks threaten the functioning of the economy itself, protecting beneficiaries may require investors to look beyond the individual security and consider the health and resilience of the larger systems in which all investments operate.

That question—how investment institutions should respond when portfolio returns depend on the stability of the system as a whole—remains an important agenda for further Sustainable Finance Forum discussion.

****

Session moderated by:
Dr. Robert K. Massie, Jr.
Director, SIRI Pathways to Consensus
Sustainable Investing Research Initiative (SIRI)
Columbia University
rkm2158@columbia.edu