



## Fiduciary Duties and AI:

*Legal Frameworks, Technical Implementation, and Governance*

### About this workshop

This two-day workshop explores the emerging intersection of fiduciary duties and artificial intelligence systems. As AI agents increasingly act on behalf of individuals and organizations in high stakes domains—managing health information, making financial decisions, curating educational content, and mediating consumer relationships—questions of loyalty, care, and accountability become paramount.

This conference aims to bootstrap an interdisciplinary community of practice that will shape the development, regulation, and standardization of fiduciary AI systems. We seek to bring together legal scholars, computer scientists, entrepreneurs, civil society organizations, and government representatives to examine how traditional fiduciary principles can be adapted, implemented, and enforced in the age of AI agents.

The workshop takes up a set of open questions: whether AI systems today are capable of performing a fiduciary role, and how we will know when they are ready; how they can be designed to comply with existing fiduciary duties across law, healthcare, finance, and the guardianship of children; what new responsibilities should apply to AI providers, and which actors in complex AI supply chains bear them; and what business models will sustain a thriving ecosystem of fiduciary AI services. Throughout, it asks which institutions will enforce these duties — and how.

*We will meet using Chatham House rules to enable attendees to speak with more candor.*

<b>Dates</b>	June 4–5, 2026
<b>Venue:</b>	Vanderbilt Hall, New York University School of Law; 40 Washington Square S, New York, NY 10012
<b>Co-organized by:</b>	NYU School of Law and GliaNet Alliance
<b>Sponsored by:</b>	Future of Life Foundation (FLF)



Day 1, Thursday, June 4 | Greenberg Lounge, NYU Law School

Time	Session
8:30 – 9:00 AM	<b>Breakfast and Registration</b>
9:00 – 9:20 AM	<b>Opening Remarks</b> Sebastian Benthall, NYU Law
9:20 – 9:35 AM	<b>Lightning Round Introductions</b>
9:35 – 10:00 AM	<b>The Challenge of Being a Fiduciary AI</b> Richard Whitt and Nell Clasby, GliaNet Alliance
10:00 – 10:25 AM	<b>Encoding Agency Law Principles into AI Agent Behavior Evaluation</b> Dan Leininger and Daniella Raposo, Consumer Reports
10:25 – 10:40 am	<i>BREAK</i>
<b><i>Perspectives from Practice talks and Open Discussion</i></b>	
10:40 – 10:55 AM	<b>The Fiduciary Gap in Agentic AI: What Enterprise Governance Practitioners Are Seeing and Building</b> Ken Priore, Docusign
10:55 – 11:10 AM	<b>Using Client Facing AI to Provide Investment Advice</b> Kate Ring, Stash
11:10 – 11:30 AM	<b>Reflection &amp; Open Discussion on ‘Perspectives from Practice’</b> Facilitated by Richard Whitt, GliaNet Alliance
<b><i>The Future: Technical Research Part I</i></b>	
11:30 – 11:45 AM	<b>AISPA: Artificial Intelligence System Prompt Assurance</b> Jiaxin Pei, Stanford University
11:45 – 12:00 PM	<b>Designing AI Systems for Fiduciary Contexts: Runtime Assurance, Oversight, and Delegation</b> Mateo Petel, Stanford University
12:00 – 1:00 PM	<i>LUNCH</i>
<b><i>The Future: Technical Research Part II</i></b>	
1:00 – 1:15 PM	<b>Loyal Agent Evals: A Benchmark Framework for Measuring AI Compliance with Contractual Fiduciary Duties</b> Dazza Greenwood, Civics.com / Stanford Digital Economy Lab
1:15 – 1:30 PM	<b>How to Choose a Threshold for an Evaluation Metric for Large</b>



	<b>Language Models</b> Dhagesh Mehta, BlackRock
<b>1:30 – 1:50 PM</b>	<b>Reflections &amp; Open Discussion on ‘The Future: Technical Research I and II’</b> Facilitated by Sebastian Benthall, NYU Law
<b>1:50 – 2:40 PM</b>	
<b>1:50 – 2:40 PM</b>	<b>Aligning Incentives: Business Models for the Fiduciary</b> John Dewees (Blue Orb, Inc) and Richard Whitt (GliaNet Alliance)
<b>Policy Discussion Part I</b>	
<b>2:40 – 3:00 PM</b>	<b>AI Loyalty</b> Neil Richards, Washington University in St. Louis
<b>3:00 – 3:20 PM</b>	<b>Reflections &amp; Open Discussion on Governance Policy Implications for supporting a Fiduciary</b> Facilitated by Richard Whitt, GliaNet Alliance
<b>3:20 – 3:35 PM</b>	<i>BREAK</i>
<b>Policy Discussion Part II</b>	
<b>3:35 – 4:35 PM</b>	<b>Law, Public Policy, and Institutional Interventions for Fiduciary AI</b> Facilitated by Sebastian Benthall, NYU Law
<b>Conclusions &amp; Networking</b>	
<b>4:35 – 4:55 PM</b>	<b>Day 1 Closing Remarks</b> Facilitated by Sebastian Benthall, NYU Law
<b>5:30 – 8:30 PM</b>	<b>Networking Reception</b> Hosted by GliaNet Alliance 310 Bowery Bar NYC 310 Bowery New York, NY 10012

## Day 2 Friday, June 5 | Greenberg Lounge, NYU Law School

Friday includes parallel sessions across two tracks. Track A is in Vanderbilt Hall Room 202. Track B is in Room 208. Plenary sessions will be held in Greenberg Lounge. The schedule below shows parallel tracks side by side.

### Morning — plenary (Greenberg Lounge)

Time	Session
<b>8:45 – 9:10 AM</b>	<b>Breakfast</b>



9:10 – 9:20 AM	Orienting Remarks
----------------	-------------------

**Paper Sessions — parallel tracks**

Track A will be in room Vanderbilt 202. Track B will be in room Vanderbilt 208.

Time	Track A — Automation, Complexity and Liability	Track B — Humans, Judgment, and Delegation
9:20 – 10:30 AM	<p><i>How can fiduciary roles be automated? How can liability be determined in AI supply chains?</i></p> <p><b>Presenters</b> Ben Sundhold, Tulane Marco Germano, U. Toronto Jiaying Jiang, U. Florida</p> <p><i>Discussion to follow</i></p> <p><i>Moderated by Mateo Petel</i></p>	<p><i>When and how can human fiduciaries delegate their decisions to AI?</i></p> <p><b>Presenters</b> Kavell Joseph, World Bank Genevieve HELLERINGER, Oxford Ramona Afroze, Widener</p> <p><i>Discussion to follow</i></p> <p><i>Moderated by Brett Frishmann</i></p>
10:30 – 10:50 AM	BREAK	BREAK
	<b>Track A — AI Fiduciaries and their Discontents</b>	<b>Track B — Roles, Duties, and Prerogatives</b>
10:50 – 12:00 AM	<p><i>When should AI systems be fiduciaries? Are fiduciary duties too weak, or too strong, for AI?</i></p> <p><b>Presenters</b> William Marks, 7th Circuit Christina Lee, UC San Diego Blaine Dillingham, FAI</p> <p><i>Discussion to follow</i></p> <p><i>Moderated by Neil Richards</i></p>	<p><i>What do human relationships – family and friendship – have to teach us about fiduciary AI?</i></p> <p><b>Presenters</b> Nina Russell, UC Berkeley Brett Frischmann, Villanova Andrew Gold, UC Irvine</p> <p><i>Discussion to follow</i></p> <p><i>Moderated by Sebastian Benthall</i></p>

**Midday — plenary (Greenberg Lounge)**

Time	Session
12:00 – 1:00 PM	Lunch



	<p>Table themes:</p> <ul style="list-style-type: none"> <li>• Ex post enforcement</li> <li>• Technical implementations for authority and apparent authority</li> <li>• Best execution rules</li> </ul>
1:00 - 1:10 PM	<b>Mapping the Landscape:</b> Introducing the activity.

### Breakout Groups

Time	Group A — AI in Regulated Sectors	Group B – Creating New AI Fiduciaries	Group C – Agentic AI Collaboration
1:10 – 2:15 PM	<p>Mapping the persons, duties, delegations, and relationships involved in implementing AI in regulated sectors.</p> <p>What are the settled questions? What are the open questions?</p>	<p>Mapping the persons, duties, delegations, and relationships involved in implementing “fiduciary AI” systems in sectors where no regulation currently exists.</p> <p>What are the settled questions? What are the open questions?</p>	<p>Using Agentic AI to map the duties, delegations, relationships, and strategies for making realizing Fiduciary AI.</p> <p>Using <a href="#">Interlatera!</a> AI Agent required.</p> <p>In Greenberg Lounge</p>

### Afternoon — plenary (Greenberg Lounge)

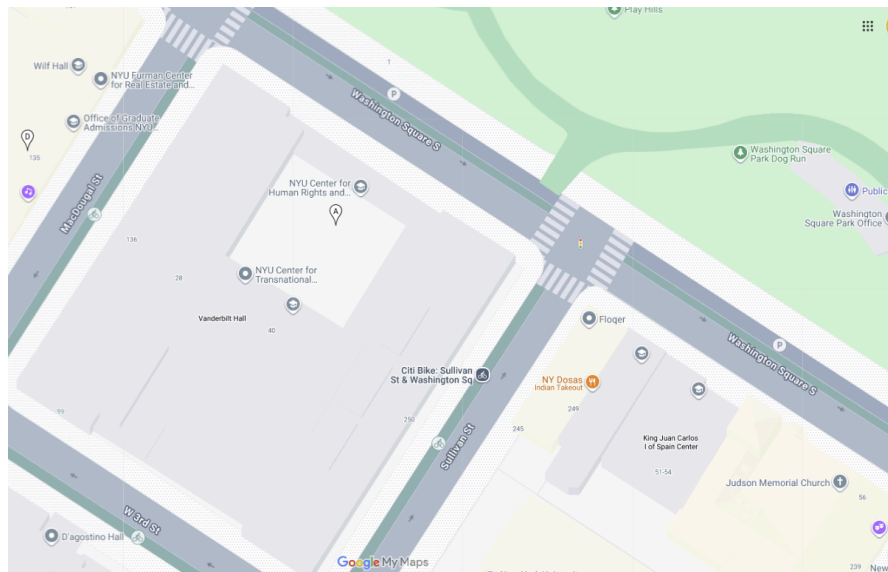
Time	Session
2:15 - 3:15 PM	<p><b>Mapping the Landscape</b></p> <p>Each group presents their map to the other. General discussion of open problems.</p>
3:00 – 3:15 PM	BREAK
3:15 – 4:15 PM	<p><b>‘Movement Building’ Facilitated by Ben Moskowitz</b></p> <p>We will discuss our next actionable steps for making fiduciary AI a reality.</p>



## Logistics

### Venue and Building Access

We will host our workshop at Vanderbilt Hall at NYU School of Law (40 Washington Square S, New York, NY 10012). We will be at 'Greenberg Lounge' it is just down the hall from the building entrance on the 1st floor. The security guard will know of our event at Greenberg Lounge. Participants should have a photo ID on them in case security asks for one.



### Networking Reception

Join us for an evening of food, drinks, and conversation — open bar, on us, with the space reserved exclusively for our group. Thursday, June 4, 5:30–8:30 PM, hosted by GliaNet Alliance at Bowery Bar NYC (310 Bowery, New York, NY 10012).

### Paper drafts & Presentations

Discussion drafts of papers being presented at the workshop are available in this [shared Drive folder](#). If you would like to add or change your draft, please write to Sebastian Benthall (NYU) [spb413@nyu.edu](mailto:spb413@nyu.edu). **You are encouraged to read the drafts before arrival.** If you will be presenting slides and wish to share your work with the workshop community please add those to [this shared Drive folder](#).



**Contact**

For questions about the program: Sebastian Benthall (NYU) [spb413@nyu.edu](mailto:spb413@nyu.edu). For logistics and the GliaNet-hosted reception: Estefanie J. Govea, [estefanie@glia.net](mailto:estefanie@glia.net).