

Murad Farzulla

Multi-Agent Systems · AI Safety · Computational Finance

London, UK · murad@dissensus.ai · dissensus.ai/research · [Google Scholar](https://scholar.google.com/citations?user=0009-0002-7164-8704) · [ORCID 0009-0002-7164-8704](https://orcid.org/0009-0002-7164-8704) · [GitHub](#)

Research Profile

I work on measurement validity in complex adaptive systems: whether a quantitative claim about a market, an institution or a population of interacting agents survives being checked. The through-line across my work is that apparent structure is frequently manufactured by the design of the study rather than present in the system, and that finding out which is a methodological problem worth taking seriously in its own right. Two of my results have died to their own controls; both were published that way.

From February 2027 I begin a mathematics PhD at City St George's, University of London, supervised by Professor Andrea Baronchelli, on the effects of delegation on coordination in multi-agent systems and what follows for alignment.

Education

PhD Mathematics

from Feb 2027

City St George's, University of London

London, UK

Supervisor: Prof. Andrea Baronchelli. Delegation, coordination and alignment in multi-agent systems. Start date agreed; formal offer letter pending.

MSc Finance Analytics

2024–2026

King's College London

London, UK

- Merit (68.25%), awarded April 2026. Distinctions in Big Data Analytics (78), Analytics in Finance, and Empirical Finance.
- Dissertation developed into an article forthcoming in *Digital Finance*.

BSc Accounting & Finance

2021–2024

SOAS, University of London

London, UK

- First Class Honours.

A-Levels: Economics, Mathematics and Computer Science

Oakham School, 2021

BTEC Level 2 Diploma — Pearson; Teamwork and Personal Development in the Community

2021

Publications

Forthcoming

Do Cryptocurrency Markets Differentiate Infrastructure from Regulatory Shocks? A Multi-Moment Event Study with Dependence-Robust Inference

GJR-GARCH-X; multi-moment event study; Student-*t*-copula bootstrap. 50 events, six cryptoassets, 2019–2025. *Digital Finance* (Springer). [Preprint](#) · [Estimator](#) · [Inference toolkit](#).

Under review

ASRI: An Aggregated Systemic Risk Index for Cryptocurrency Markets (*with A. Maksakov*)

Multi-source pipeline, 1,841 observations; PCA; HMM regime detection; walk-forward evaluation; moving-block bootstrap; public researcher API. *Data-Driven Modelling*. [arXiv:2602.03874](#) · [Live dashboard](#).

Are Whitepaper Claims Reflected in Market Structure? A Contamination-Aware Pipeline and a Power-Limited Null

Dataset reconstruction; NLP classification; tensor decomposition; contamination controls. *Financial Innovation* (SpringerOpen). [arXiv:2601.20336](#).

From Consent to Consideration: Functional Criteria for AI Political Standing

Four functional criteria for when governing an autonomous system is legitimate. *AI and Ethics* (Springer). [Zenodo](#).

Relational Functionalism: A Functional Account of Substrate-Independent Friendship

Relationships as substrate-agnostic functional processes, implementable across biological and artificial systems. *Ethics and Information Technology* (Springer). [Zenodo](#) · [PhilArchive](#).

Preprints

Four Ways to Manufacture a Coordination Finding: A Control Battery for Mixed-Motive MARL

IQL and VDN; 125-condition factorial, 30 seeds, 3,750 runs; approximately 15,000 runs across the full experimental programme. Separates structural effects from reward scaling, preference construction and sampling artefacts. [Research Square](#) · [Dataset](#) · [Code](#). Live preprint carries a superseded title; retitled version pending.

The Replicator-Optimization Mechanism: A Scale-Relative Formalism for Persistence-Conditioned Dynamics

Persistence-conditioned dynamics with application to consent-based metaethics. [arXiv:2601.06363](#) · [Zenodo](#).

Axiom of Consent: Friction Dynamics in Multi-Agent Coordination

Formal treatment of coordination friction; companion to the MARL control battery. [arXiv:2601.06692](#) · [Zenodo](#).

The Extremity Premium: Sentiment Regimes and Adverse Selection in Cryptocurrency Markets

Agent-based model; sentiment intensity rather than direction drives uncertainty-linked liquidity withdrawal; 2,896 trading days. [arXiv:2602.07018](#) · [Zenodo](#).

Privacy-Preserving Financial Surveillance: An Architectural Framework for CBDC Implementation *(with A. Maksakov)*

Zero-knowledge proof architecture; private set intersection for sanctions screening; the trade-off between transaction privacy and supervisory visibility. [Zenodo](#) (concept DOI; v3.0.0, corrected August 2026). A retitled and restructured version is in preparation; the edited-volume citation below refers to this title.

Context-Dependent Affordance Computation in Vision-Language Models

Systematic cross-model evaluation (Qwen3-VL, LLaVA); latent-factor stability; replication across model families. [arXiv:2603.04419](#) · [Zenodo](#).

Trauma as a Functional Kind: Multiple Realizability and the Individuation of Maladaptive Learning

Machine learning and developmental psychopathology; four PyTorch models with a reproducible figure pipeline. [Research Square](#) · [Code](#).

Alpha Asymmetry in Foreign Exchange Markets: An Investigation of Exploitability

Extreme-value methods applied to cross-market efficiency; a null result. [Zenodo](#) · SSRN:6147567. Corrected version pending deployment.

Legitimate Extraction: Sophisticated Laundering Hides in Plain Sight

A fourth stage of money laundering; bears on financial-crime detection and AML enforcement design. SSRN:6145046 · [Zenodo](#).

Consent-Theoretic Framework for Quantifying Legitimacy: Stakes, Voice, and Friction in Adversarial Governance

Legitimacy quantification in adversarial institutional settings. [Zenodo](#) · SSRN:5918222.

Autonomous Red Team: Multi-Agent Adversarial Security Testing *(with A. Maksakov)*

LLM-guided adversarial security testing; red/blue publish-gate evaluation design. [Zenodo](#).

Stakes Without Voice · Identity is Irreducibly Relational · Genre Mimicry in Academic Writing · Dissolving Qualia via Occam’s Razor (monograph, 33k words) · The Trident · Quantum Statistics from Oscillatory Sampling

Further preprints in governance, philosophy of mind, philosophy of mathematics and foundations of physics. Full archive with DOIs: dissensus.ai/research.

External Citations

Independent citations only; self-citations and citations in venues I judge predatory are excluded from this list, and Google Scholar’s count is therefore higher than what appears here. Law-review and edited-volume citations are invisible to Scholar, OpenAlex and Semantic Scholar, and were found manually.

- **ASRI** — Uri Volovelsky (Israeli Ministry of Justice) & Sivan Shlomo Agon (Bar-Ilan), “Decentralized Autonomous Organizations and the Anti-Money Laundering Challenge: Rethinking Global Frameworks for a Leaderless World”, *17 Harvard National Security Journal 2 (2026)*, notes 374–375. ASRI is the anchor source for their proposed DAO Early-Warning Index, a core recommendation of the article — uptake by legal practitioners rather than academic citation alone.
- **The Extremity Premium** — Huang & Zhang, “Act or hesitate? Investor sentiment divergence and informed trading”, *Pacific-Basin Finance Journal* (Elsevier, 2026).
- **Relational Functionalism** — Jaime Banks (Katchmar-Wilhelm Professor, Syracuse iSchool), “Ghosting the Machine: Stop Calling Human-Agent Relations Parasocial”, [arXiv:2604.05197](#) (2026), forthcoming in ACM proceedings.
- **ASRI** — David Krause (Marquette University), “Fragile Links: Private Credit Tokenization and DeFi Contagion”, SSRN (2026).
- **ASRI** — Oberholzer & Zamaraiev, “Toward a Risk Assessment Framework for Institutional DeFi: A Nine-Dimension Approach”, [arXiv:2605.05145](#) (2026).
- **ASRI** — Tamer Abdelaziz (New York University), “Cross-Chain Exploit Propagation and Systemic Risk Modeling”, working paper (2026).
- **ASRI** — Zhengpu Zhao (UC Berkeley), “Decentralization, Crowds, and Cycles: Why Bitcoin’s Architecture Makes Booms and Busts Structurally Likely”, SSRN (2026).
- **Whitepaper Claims** — S. Han, “Price Formation, Order Book Integrity, and Execution Continuity in Virtual Asset Markets”, SSRN (2026).
- **CBDC Privacy Architecture** — cited in *Intelligent Embedded Systems 2: Advanced Architectures, IoT Integration and Autonomous Decision-Making*, eds. Kumar, Batta, Ravi & Mohan, ISTE Ltd / Wiley, Computer Engineering Series (2026).

Ongoing Research

Population-level evaluation for LLM multi-agent systems

2026–

Primary research programme

- Reproducing an established LLM-population coordination baseline, then running a control battery that separates individual predisposition, information and memory conditions, and interaction structure — asking which population-level effects survive controls at all. Direct successor to the MARL result below, and the empirical groundwork for the doctorate.

ASRI — live systemic-risk index

2025–

With A. Maksakov

- Production monitoring system regenerating daily from committed code with integrity guards, channel-decomposed computation, a frozen data record and a public researcher API on Cloudflare Workers. Live at asri.dissensus.ai.

Formal verification of the theoretical corpus (Lean 4)

2026–

- Machine-checked mechanisation of my theoretical papers: 52 files, 246 theorems, zero `sorry` and zero custom axioms, building clean against Mathlib v4.27.0. dissensus-ai/lean-formalizations

Compute and research resources

2026

- Dual-GPU workstation (7900 XTX + 7800 XT, 128 GB) running open-weight model inference and training locally under ROCm — the full seed grid on one model family costs nothing marginal.
- Apodex Frontier Program: \$500 in API credits awarded to Dissensus Ltd, 2026.

Open-source research software

2025–

Maintainer

- `gjr-garch-x` (GJR-GARCH-X volatility estimation) and `robust-eventstudy` (dependence-robust event-study inference) on PyPI — continuous integration, golden-file test suites, ODC trusted publishing, Zenodo archives.
- Multi-agent evaluation harnesses that verify papers against their committed code and adversarially cross-examine review findings; red/blue adversarial publish-gate design.

Collaborations

- **Andrew Maksakov** — Research Assistant, Dissensus. Co-author on ASRI, the CBDC privacy architecture and Autonomous Red Team.
- **Prof. Andrea Baronchelli and Ariel Flint Ashery** (City St George's, University of London) — collaboration on conventions and collective bias in LLM populations, ahead of the February 2027 doctorate.
- **Dr Ryan Quandt** (Senior Economist, Council of Economic Advisers; previously Chief Economist and Director of Research, Cicero Institute — corresponding in a personal capacity) — proposed empirical programme testing the Agonism Hypothesis for human–AI teaming: whether structured dissent improves decision quality, and where its marginal yield stops paying. I contributed a Lean 4 formalisation establishing that the optimal dissent dose is invariant to stakes and alignment while the penalty for mis-dosing is not. Under discussion.
- **Matthew Cornelius** (*Philosophical Psychology*, 2026) — joint paper in draft on two independent routes to illusionism about consciousness. His architectural “wet room” account explains the character of the seeming but leaves the presentation layer unmotivated; my replication-optimization mechanism explains why selection would build one but is silent on implementation. Each account is incomplete where the other is not. First section drafted August 2026.
- **Romina Roca** — joint paper, *Operationalizing the Noosymbiotic-Extractive Demarcation: A Dependency-Slope Approach to Human-AI Relational Trajectories*. The noosymbiosis framework is hers; the dependency-slope operationalisation is my contribution.
- **Dr Felipe Pachano Azuaje** (Sineidocracia Foundation) — editorial board member; contributing author on two chapters of the Foundation's volume, and reviewer on a third.
- **Computational Philosophy Lab**, Northeastern University London — affiliate, working with Prof. Brian Ball and Prof. David Freeborn; agent-based modelling and collective behaviour.

Experience

Independent Researcher & Founder

2025–present

Dissensus Ltd

London, UK

- Lead empirical projects in computational finance and multi-agent systems end to end: research design, data acquisition, analysis, validation, revision and reproducible release.
- Coordinate a research assistant and external collaborators across three countries; maintain documented workflows and translate findings into papers, technical documentation and public dashboards.

Analysis Editor, Eurasia

2026–present

Third Institute

Remote

- Commission and edit policy and economic analysis for a non-specialist readership; test claim–source alignment, improve structure and clarity, coordinate revisions to deadline.

Operations Manager & Project Consultant

2023–2025

Prince MMC

London / Baku

- Financial forecasts, supplier and contract records, operational reports; delivery coordination across commercial and public-sector engagements. Clients included Azerbaijan Airlines and government bodies.

Training, Service and Affiliations

Technical AI Safety and Frontier AI Governance — BlueDot Impact, professional certifications awarded July 2026 (competitive cohorts).

Introduction to Cooperative AI — Cooperative AI Foundation, 29 June–28 August 2026, by application. Multi-agent cooperation, commitment, coordination and conflict.

AI Agents and the Law — Vista Institute for AI Policy, 19 August–2 October 2026 (two modules, by application). How law applies to increasingly capable AI agents; includes a research component.

Digital Minds Research Sprint — Apart Research, 14–16 August 2026. Team research sprint on the empirical study of possible digital minds; tracks spanning preference elicitation, introspection and self-report, and assistant persona.

Financial Trading Programme — TTC in partnership with King’s College London, July 2025. Certificate awarded on completion.

Focus on Peer Review — Nature Masterclasses, August 2026.

Invited Peer Reviewer, *Financial Innovation* (Springer), 2026.

Affiliate, Computational Philosophy Lab, Northeastern University London, 2026–.

Editorial board member, Sineidocracia Foundation, 2026–.

Founder & Director, Dissensus Ltd (incorporated England & Wales, 2026).

Languages: English (native), Azerbaijani and Russian (fluent), Turkish (working), French (basic).

Volunteering & Community

- Volunteer coding teacher for primary school children.
- Fundraising for UK food banks and for orphanages in Azerbaijan; provision of daily meals for a children’s education centre.
- Advocacy and awareness work on neurodivergence and mental health.
- School Prefect and Combined Cadet Force (Lance Corporal), Oakham School.

Technical Skills

Programming: Python (daily — PyTorch, pandas, NumPy, SciPy, statsmodels, scikit-learn), Lean 4, R, SQL, Bash, Git/CI, L^AT_EX. **Methods:** multi-agent reinforcement learning, agent-based modelling, LLM and vision-language model evaluation, GARCH-family econometrics, tensor decomposition, extreme-value methods, bootstrap and dependence-robust inference, hidden Markov models, formal verification. **Infrastructure:** GPU clusters (ROCm/CUDA), local LLM inference and training on a dual-GPU workstation, Cloudflare Workers, Docker, Kubernetes.