Input Output Research

Cardano Vision & Work Program 2025

Mid-year Report - Fundamental Research

Executive Summary

Cardano Vision launched in 2025 with Work Program 2025 (WP25), a five-year research initiative designed to secure Cardano's long-term leadership in blockchain. Enabled by community support and grounded in peer-reviewed science, it tackles the hard challenges—sustainability, scalability, interoperability, and security. This Mid-Year Report provides a transparent update on progress, milestones, and priorities for this first year.

Evidence-Based methodology

Input | Output Research (IOR) applies an Evidence-Based methodology to ensure correctness, security, and reliability in blockchain systems—where exhaustive testing alone is insufficient. Rooted in peer-reviewed science and formal specification, each research stream progresses from well-defined problems through models, proofs, and executable artifacts, aligning theory with real-world implementation. An agile, iterative process enables early issue detection, efficient resource use, and continuous refinement. Progress is tracked through Software Readiness Levels (SRLs):

- **Fundamental Research (SRL 1–2):** formalizing ideas and proofs, 3–5 years to market
- **Technology Validation (SRL 3–5):** prototyping, validating and specs, 1.5–3 years
- **Targeted Implementation (SRL 5+):** guiding engineering into production, 0–18 months

This structured funnel supports a steady pipeline of innovation, with more than 100 research outputs expected over five years, with 30 advancing to validation and implementation. Building on Cardano's proven record, it provides the assurance needed for secure, scalable, and future-ready capabilities.

Fundamental research

In 2025, IOR is advancing 20 foundational research streams under Cardano Vision, already delivering promising outputs that reinforce Cardano's position as a leader in blockchain science. These streams formalize new ideas beyond the current state of the art, combining clear problem definitions, formal modeling, and rigorous security proofs to produce high-assurance solutions.

Outputs typically include technical reports, peer-reviewed papers, and formal specifications – critical evidence for later prototyping and validation. Recognizing the exploratory nature of early-stage research, IOR applies a flexible portfolio approach. This report includes three change requests reflecting updated stream prioritizations to ensure resources are focused where impact is greatest.

Cardano Vision

Cardano Vision is a five-year research agenda focused on sustainability, scalability, and interoperability—recognizing that transformative blockchain technologies require long-term, science-led investment rather than short-term market demand. By advancing high-potential areas today, Cardano ensures continued innovation, global competitiveness, and broad societal benefit.

Building on Cardano's record of 100% uptime, the program is structured around nine thematic focus areas. At its core, Cardano Vision seeks to evolve blockchain into a global operating system for decentralized compute, storage, and identity, enabling seamless interoperability and supporting new digital institutions for identity, governance, and trust.

Impact and outputs

The first year of Cardano Vision is progressing strongly, with a balanced portfolio of long-running and newly launched research streams—some already delivering outputs and others advancing cutting-edge directions—supported by three prioritization change requests, and with IOR on track to exceed its contractual target of over 20 research outputs in 2025 (see Appendix B).

World's Operating System (WOS): In 2025, *WOS-4* advanced both a formal treatment of adaptively secure decentralized storage networks (DSNs), and Byzantine-resilient primitives for DHTs and DAS, exploring PoSpace/PoStake Sybil defenses and potential DAS–RDA unification. *WOS-6* introduced a novel mechanism incentivizing geographic diversity, with experimental validation and work on integrating verified location claims into cryptographic protocols.

Ouroboros Omega (OO): Major progress includes *OO-1V*, a forthcoming Peras paper proving safety, liveness, and self-healing, supported by engineering reports confirming integration feasibility. *OO-2* Leios achieved a key milestone with a Crypto '25 paper on throughput scalability, while follow-up work addresses concurrency and conflicts. Additional outputs span a *OO-3* paper being presented at Asiacrypt '25 on adaptive security for fair transaction processing, *OO-5* multi-resource consensus with Ethereum/Bitcoin re-staking, *OO-6* new consensus design and early Proof-of-Deep-Learning results, and *OO-7* publications on congestion fees, space tokenization, and pricing models.

Tokenomicon (TO): *TO-1* advanced economic models for Cardano (AFT '25), extending earlier work to reserve policies and resilience, complemented by large-scale surveys of 11,000 participants and new frameworks for algorithmic monetary policies. *TO-2* introduced a Shapley-value–based pooling scheme (AFT '25), explored pooling trade-offs, and designed incentives for Mithril, with results disseminated at major conferences.

Democracy 4.0 (D4): *D4-1* finalized a UC security framework for L2 governance and designed a minimal-footprint protocol. *D4-2* modeled DRep incentives, and explored participatory budgeting for Catalyst, with results presented at FC '25 and ALGA '25.

Internet Hydra-ted (IHT): *IHT-1* developed a rollup-style protocol with design complete and a security proof due in 2025. *IHT-2* delivered the first UC framework for multiparty state channels, with extensions for channel composition in progress. *IHT-3* mapped mechanisms like fund rebalancing and routing, producing a roadmap for Hydra enhancements in 2026.

Interchains (IC): *IC-1* delivered the first formal bridge security framework, alongside ZK advances on efficient Merkle openings (ICCCN '25) and ongoing TEE bridge work. *IC-3* introduced *Cavefish*, an ultra-efficient light client for UTxO chains. *IC-4.1* advanced tokenomics for system launches (IJCAI '25) with behavioral insights of investors vs gamblers (CRETE '25). IC-4.2 progressed Jolteon formal verification (safety complete, liveness underway), developed a game-theoretic model of PoS incentives, and an anti-grinding paper (AsiaCrypt '25).

Core Zero-Knowledge (ZK): *ZK-1* advanced foundational UC frameworks, with AGATE (CSF '25) formalizing TEEs, UC-SNARKs with transparent setup (Crypto '25) and a follow up paper on TEE-based ZK proof servers. Applied work improved Mithril via succinct aggregation (BLS, Merkle proofs), while post-quantum ZK efforts produced an SoK survey on lattice-based SNARKs and folding schemes to guide future research.

Communication and dissemination

IOR actively shares research to ensure transparency, collaboration, and impact across the Cardano ecosystem:

IOHK channels: The IOHK website, YouTube, blog and social media provide a comprehensive library of papers, videos, and thought leadership. Notable posts covered consensus evolution, smart contract verification, cryptographic tools, Leios, airdrop games, SPO research frameworks, and Peras.

Cardano R&D Sessions: IOR relaunched its community engagement in June with monthly thematic discussions featuring invited guest speakers from the community. Highlights include *Layer 2 Expansion* – Beyond Hydra (June) on L2 protocol research, rollups, ZK bridges, and interoperability; Cardano Tokenomics (July) on incentives, pledge mechanics, stablecoins, and CIP-50 Rebirth; and Technology Validation (August) featuring Phalanx, Jolteon, RSnarks, Minotaur, Cavefish, and Committee Proofs.

Research conferences: In the first half of 2025, IOR presented research at major venues including FC '25, FMBC '25, CSF '25, IJCAI '25, ITC '25, Crypto '25, and others, contributing to blockchain science and advancing Cardano Vision's goals.

Contents

1. Introduction

Work Program 2025 (WP25) marks the beginning of **Cardano Vision**—a bold, five-year research initiative to shape the future of the Cardano ecosystem. This initiative would not exist without the continued trust, engagement, and support of the Cardano community. Your votes, your feedback, and your belief in science-led research and development have made WP25 possible. Input | Output Research (IOR) is grateful for your support and to be building this future with you.

Cardano Vision is grounded in a simple idea: to lead in blockchain for the next decade, we need to invest in science today. It's about solving the hard problems—scalability, interoperability, governance, privacy, identity, and security—with rigour, creativity, and purpose. The WP25 proposal sets the foundation, delivering over 20 research and 6 technology validation workstreams focused on unlocking the next wave of decentralized infrastructure.

This Mid-Year Report provides a transparent overview of IOR's fundamental research progress in 2025, outlining milestones achieved, early findings, and upcoming priorities. While aligned with the Technology Validation mid-year report, it also stands alone as a clear account of our contributions. Despite challenges from working at risk and onboarding new resources, we are pleased to present both the contractual deliverables achieved and the broader role IOR plays in advancing Cardano through world-class research and community engagement.

2. Evidence-Based methodology

Achieving correctness, security, and reliability in decentralized systems is inherently complex. Blockchains operate on global, public infrastructure and rely on volunteer node operators. Since it's impossible to test every possible state or failure mode experimentally, formal methods—mathematical models and proofs—are required to build confidence in the behaviour of these systems under all conditions.

Input | Output's methodology is rooted in peer-reviewed science and formal specification. Each research stream begins with a well-defined problem and progresses through conceptual modelling, security proofs, and eventually executable artifacts. This ensures a tight alignment between theoretical design and real-world implementation, resulting in high-assurance systems that can be confidently adopted and evolved over time.

While our assurance standards are rigorous, IOR follows an agile, iterative process. This flexible approach enables early identification of roadblocks, efficient resource use, and continuous refinement of artifacts and specifications throughout a workstream's lifecycle. Progress is guided by Software Readiness Levels (SRLs)—a structured framework inspired by Technology Readiness Levels (TRLs) that helps assess maturity, estimate time-to-market, and manage risk.

- Fundamental Research (SRL 1–2): IOR develops and formalizes ideas that push beyond the current state of the art—defining clear requirements, identifying fundamental trade-offs and limitations, constructing meaningful mathematical models, establishing well-scoped design goals, and proposing technically sound solutions supported by rigorous security proofs (3 - 5 years to market).

- Technology Validation (SRL 3–5): An interdisciplinary innovation team conducts targeted validation of early-stage concepts—developing promising ideas through rigorous R&D, demonstrating feasibility via prototypes, models, and simulations, and producing specifications that inform and support full implementation (1.5 - 3 years to market).

- Targeted Implementation (SRL 5+): IOR supports engineering teams in implementing solutions within a target production environment, guided by specifications developed during the innovation phase. This ensures an evidence-based engineering process with a strong focus on assurance and reliability (0 - 18 months to market).

This structured innovation funnel guides research from early exploration to deployment readiness. Over five years, it is expected to generate more than 100 high-quality research outputs, with around 30 advancing to technology validation and implementation. IOR has been catalytic in Cardano's growth into a multi-billion-dollar ecosystem, and this multi-year approach ensures a steady pipeline of globally relevant innovations, and accelerates Cardano's delivery of secure, scalable, future-ready capabilities that address the most pressing challenges in blockchain.

3. Fundamental research

Under WP25, IOR is advancing 20 foundational research workstreams for the first year of Cardano Vision. These streams are already producing promising outputs and highlight IOR's commitment to high-impact, long-horizon research that strengthens Cardano's role as a leader in blockchain innovation.

At their core, these workstreams aim to formalize new ideas that go beyond the current state of the art. This involves defining precise requirements, identifying trade-offs and limitations, developing mathematical models, setting clear design goals, and proposing technically sound solutions backed by rigorous security proofs.

- Problem Statement – Each workstream begins with a carefully scoped problem definition, capturing both functional and non-functional requirements. This ensures clarity of goals and constraints from the outset.

- Abstraction and Formal Modeling – Problems are then abstracted into formal models that capture only the essential aspects, with assumptions and boundaries documented. This step assesses feasibility and guides refinements if a given path proves unviable.

- Solution Design and Security Proofs – Within the formal model, candidate solutions are developed and their properties—such as security, correctness, or liveness—are mathematically proven under well-defined assumptions.

- Documentation and Peer Review – Outputs are compiled into structured reports and peer-reviewed publications. Where appropriate, results are also shared through Cardano Improvement Proposals (CIPs) or Intersect working groups, helping connect theoretical advances to implementation pathways.

The outputs of this process typically take the form of technical reports, draft or final research papers, and in some cases formal specifications. These artifacts are critical stepping stones, providing the evidence base needed to move research into prototyping and validation phases.

Foundational research is inherently uncertain and exploratory. Priorities may shift as new insights emerge and hypotheses are refined, as is expected in any rigorous academic process. For this reason, a flexible, adaptive portfolio approach is essential—ensuring Cardano remains responsive to emerging opportunities while continuing to expand its research frontier.

4. Cardano Vision

Guided by a five-year Strategic Research Agenda focused on sustainability, scalability, and interoperability, Cardano Vision recognizes that deep technologies cannot rely solely on market demand. They require sustained, early-stage research investment. By advancing high-potential areas now, Cardano accelerates innovation, secures long-term competitiveness, and delivers broad economic and societal benefits.

Building on Cardano's proven track record of 100% uptime, the program is structured around nine thematic focus areas, each addressing a critical bottleneck or opportunity. Together, these streams expand Cardano's world-class research portfolio, reinforce its standing in the global scientific community, and provide a steady pipeline of commercialization opportunities for product teams and ecosystem builders.

I. **The World's Operating System (WOS):** To support decentralized applications and digital assets at scale, Cardano must offer powerful, secure, and feature-rich smart contract capabilities. This enables value transfer and more complex interactions, such as DeFi, while maintaining performance and security.

II. **Ouroboros Omega (OO):** As Cardano's proof-of-stake consensus protocol, Ouroboros provides energy efficiency and resilience. However, continued research is needed to meet rising demands for throughput and performance as the network scales.

III. **Tokenomicon (TO):** Tokenization powers value representation on blockchain. Cardano leads with features like native assets and Babel fees, but the economic models that underpin token ecosystems remain underexplored. This stream focuses on rigorous tokenomics research and design.

IV. **Global Identity (GI):** Decentralized identity empowers users with privacy-preserving, user-centric, and interoperable digital credentials. Research here aims to define the foundations of a global identity system and integrate it within the Cardano ecosystem.

V. **Democracy 4.0 (D4):** Building on historical evolutions of democracy, this stream explores how blockchain can support scalable, inclusive, and secure governance mechanisms, adapting democratic processes to meet 21st-century challenges.

VI. **The Internet Hydra-ted (IHT):** Hydra is Cardano's state-channel suite for layer 2 scalability. While the Hydra Head is deployed, full functionality requires continued R&D to enable high-performance, low-latency applications comparable to Web2 platforms.

VII. **Interchains (IC):** Interoperability across blockchains and external systems is critical but comes with major security challenges. This stream develops evidence-based, minimal-trust mechanisms to securely connect Cardano to other ecosystems.

VIII. **Core Zero-Knowledge Capabilities (ZK):** Zero-knowledge proofs (ZKPs) are key to privacy and efficiency. Cardano needs modular, updatable ZK tooling that supports applications like light clients, state proofs, and bridges—keeping pace with rapid advances in ZK research.

IX. **The Post-Quantum Landscape (PQL):** Quantum computing threatens today's cryptographic foundations. This stream focuses on replacing vulnerable primitives and designing quantum-secure consensus and key management systems—while also exploring quantum-enhanced blockchain capabilities.

At its core, Cardano Vision envisions blockchain evolving into a global operating system for decentralized compute, storage, and identity—where networks interoperate seamlessly, much like today's internet. This vision supports the emergence of new digital institutions, redefining how identity, governance, and systems of trust function in a connected world.

For a full list of work streams please visit the Input Output Research (IOR): Cardano Vision - Work Program 2025 full proposal on Ekklesia (see Supporting Links section).

5. Impact & outputs

The first year of Cardano Vision is well underway, with around ten streams building on prior work from IOR's long-running Blockchain Research program, each representing areas of sustained investigation over multiple years. The remaining fifteen or so streams are more recent, with roughly half initiated in the last 12-18 months in anticipation of Cardano Vision and the others commencing in 2025. This report also includes three change requests with *WOS-4 Decentralized Storage*, *IC-1 State Proofs & Bridges*, and *ZK-1 Core zero-knowledge capabilities prioritized as* outlined in Section 7.

Together, this creates a balanced portfolio—some streams are already delivering tangible research outputs, while others push into cutting-edge directions that will yield impact over the next one to two years. In total, IOR is pleased to report it is on track to comfortably exceed the research outputs it is contractually required with over 20 papers either published, accepted, under submission or manuscripts on schedule for completion this year.

World's Operating System (WOS)

In 2025 two streams were advanced; *WOS-4 Decentralized storage* *(change request)* developed Byzantine-resilient primitives for DHTs and DAS, exploring proof-of-space and proof-of-stake for Sybil resistance and assessing whether unifying DAS with Robust Distributed Arrays could improve efficiency. *WOS-6 Location-based services and smart contracts* introduced a novel mechanism to incentivize geographic diversity in decentralized networks, presented in *"Incentivizing Geographic Diversity for Decentralized Systems"* (under submission), with experimental validation and ongoing work on integrating verified location claims into cryptographic protocols.

Ouroboros Omega (OO)

Under Ouroboros Omega major advances have been achieved in 2025 to date. *OO-1V Peras* progressed toward a fast-settlement protocol with a forthcoming paper proving safety, liveness, and self-healing, alongside engineering reports showing integration atop Praos is feasible with minimal disruption. Work also optimized committees, certificate storage, and communication, while interim research evolved into the Anti-Grinding workstream. *OO-2 Leios* delivered a key milestone with a Crypto '25 paper demonstrating throughput scalability under realistic assumptions, with follow-up work addressing concurrency and transaction conflicts in parallel block production.

Other streams advanced fairness, resource use, and efficiency. OO-3 Fair transaction processing produced multiple top-tier papers formalizing transaction ordering with bounded unfairness. OO-5 Minotaur began developing multi-resource consensus with Ethereum and Bitcoin re-staking, exploring deterministic contributor-based committees and liveness recovery. OO-6 Proofs of useful work extended Ofelimos with a practical consensus design (paper under submission) and demonstrated early Proof-of-Deep-Learning feasibility on benchmark tasks. OO-7 Congestion control advanced predictable service and multi-resource fee models, with published results on fee mechanisms and space tokenization, a submission on one- vs multi-dimensional pricing, and a thesis proposing derivative-based service guarantees.

Tokenomicon (TO)

In 2025, the Tokenomics focus area has advanced both theory and practice. TO-1 Tokenomics design developed new equilibrium models for Cardano's economy, extending work from *Single-token vs Two-token Blockchain Tokenomics* (AFT '25) to cover reserve policies, adoption dynamics, and resilience under shocks. Large-scale surveys of 11,000 participants across four countries provided first-of-its-kind behavioral data, now being integrated into refined models, alongside new frameworks for algorithmic monetary policies.

TO-2 Rewards sharing and transaction fees introduced a Shapley-value–based pooling scheme (*Pool Formation in Oceanic Games*, AFT '25) as an alternative to proportional sharing, compared on-chain vs off-chain pooling trade-offs, and designed incentives for applications like Mithril certificate signing. Results were widely disseminated through conferences and workshops, reinforcing Cardano's leadership in blockchain economic design.

Democracy (D4)

In 2025 the focus area advanced governance design and incentives. D4-1 Next-level governance protocols finalized a UC security framework for layer-2 governance and, using it, designed a protocol with strong guarantees and minimal footprint, detailed in *Beyond Blockchain Ballots* (under submission). D4-2 Governance incentives published *Reward Schemes and Committee Sizes in Proof of Stake Governance* (FC '25), modeling how rewards, delegation, and effort shape DRep outcomes, offering guidance for Voltaire. Additional work explored participatory budgeting models relevant to Catalyst, with results presented at FC '25 and ALGA '25.

Internet Hydra-ted (IHT)

IHT-1 Hydra Tail developed a rollup-style multi-party state channel protocol with design and high-level code largely complete, and a research paper with full security proof planned for 2025. *IHT-2 Inter-Head* delivered the first Universally Composable (UC) framework for multiparty isomorphic state channels, proving Hydra's security and submitting results in *"Universally Composable Treatment of Multi-Party Isomorphic State Channels."* Preparatory work began on extending this to channel composition, with proofs targeted for 2026. *IHT-3 Optimization Tools* surveyed techniques like fund rebalancing, routing, and synchronization, producing a roadmap for adapting them to Hydra in 2026. Collectively, these advances strengthen Hydra's security, scalability, and efficiency.

Interchains (IC)

The Interchains focus area delivered key advances across bridges, light clients, tokenomics, and consensus. *IC-1 State proofs and blockchain bridges (change request)* produced *Bridge Games* (under submission), the first formal security framework for cross-chain bridges, alongside ZK research on efficient Merkle openings (ICCCN '25) and ongoing TEE bridge work. *IC-3 Light Client Infrastructure* introduced *Cavefish* (under submission), a highly efficient light client protocol for UTxO blockchains requiring only two rounds of interaction.

IC-4.1 DApps tokenomics advanced "tokenomics for system launches," modeling equilibria in airdrop-based participation (*Airdrop Games*, IJCAI '25; *Self-Reinforcing Airdrops*, under submission) and integrating behavioral insights from surveys (*Investors vs Gamblers*, CRETE '25). *IC-4.2 Consensus innovation* progressed formal verification of Jolteon (safety proofs complete, liveness underway) and developed a new game-theoretic model explaining PoS block-minting incentives. Together, these results lay the foundations for secure, efficient, and economically robust cross-chain interoperability.

Core Zero-Knowledge Capabilities (ZK)

Under *ZK-1 Core zero-knowledge capabilities (change request)*, *AGATE* introduced a UC framework for TEEs (CSF '25), while *UC-Secure Zero-Overhead SNARKs* (TCC '24) proved Groth16 can achieve UC security without efficiency loss, strengthening the foundations for Cardano's ZK infrastructure. Applied work enhanced Mithril with more succinct aggregation via BLS signatures and Merkle proofs, while early post-quantum research on lattice-based SNARKs and folding schemes produced an SoK survey to guide future directions.

6. Research stream updates

For each research stream, we provide an overview of its objectives, current status against the work plan, key deliverables, and results. Research outputs are typically peer-reviewed papers submitted to top-tier conferences, supported where relevant by simulations and executable artifacts to guide future prototyping and engineering. Performing peer-reviewed research is central to establishing and maintaining excellence across all streams. The schedule of publishable units is continually reviewed and refined as streams progress, with any adjustments documented through formal change requests (see Section 7).

WOS-6: Location-Based Services and Smart Contracts

Start date: Oct 22
Duration (months): 48

(i) Stream overview

The Location-based Services and Smart Contracts workstream explores how geographic location can be integrated into blockchain protocols and smart contract execution. By enabling contracts to respond to physical location data, it opens up novel applications such as location-based payments, region-specific DeFi rules, and automated billing for physical infrastructure access.

In parallel, the project investigates how node location affects network resilience. Greater geographic diversity among consensus nodes reduces exposure to geopolitical interference, natural disasters, and eclipse attacks. To support this, the workstream is developing a location verification protocol and exploring incentive mechanisms that reward operators for globally distributed node placement. Together, these advances strengthen Cardano's decentralization, expand its DApp possibilities, and align blockchain infrastructure with real-world geographic dynamics.

(ii) Context and objectives

Current blockchain systems lack robust methods for verifying the physical distribution of nodes or embedding location into contract logic. Previous research, such as the VerLoc protocol, has shown how verifiable localization can be achieved in decentralized systems. Building on this, the workstream addresses two interconnected challenges:

1. **Secure Location-Based Smart Contracts** – Design contracts that incorporate location as a condition for execution, enabling use cases like access control, compliance with regional regulations, or payments tied to physical presence.
2. **Incentivizing Geographic Diversity** – Develop a location verification protocol that allows stake pool operators (SPOs) to prove their locations with cryptographic assurance, and integrate this into a reward-sharing scheme. Nodes in underrepresented or remote regions would receive higher rewards, strengthening global resilience.

Minimal requirements are a location verification primitive that works under adversarial conditions with predefined accuracy, and a modified reward system that uses this information to rebalance incentives. By meeting these objectives, Cardano can ensure its network is both technically robust and geographically inclusive, while unlocking a new class of location-aware decentralized applications.

(iii) Work performed and main achievements

In the first half of 2025, this stream focused on the formalization and theoretical analysis of distributed location verification. The main achievement was the development of a novel mechanism for incentivizing geographic diversity in permissionless networks, ensuring nodes truthfully report and diversify their locations.

This work, presented in *Incentivizing Geographic Diversity for Decentralized Systems*, *Raghav Bhaskar, Aggelos Kiayias, Evangelos Markakis, Marc Roeschlin* (under submission), introduces a graph-theoretic framework that uses round-trip time (RTT) measurements to detect location spoofing and applies game-theoretic analysis to prove that honest location reporting forms an equilibrium. Experimental validation further quantified adversarial success rates and demonstrated the practicality of the approach. Alongside this publication effort, discussions have begun on how verified location claims—particularly from SPOs—can be integrated into cryptographic protocols such as key generation, setting the stage for applied follow-up work.

Deliverables

#	Description	Status
D1	An initial paper covering the location verification primitive and protocol: *Incentivizing Geographic Diversity for Decentralized Systems*	Under submission '25
D2	Paper on geographic diversity	In progress

(iv) Results beyond state of the art

This research represents the first rigorous treatment of geolocation in distributed, fully untrusted systems. By combining secure RTT-based localization with incentive-compatible mechanisms, it advances beyond prior work on verifiable localization, which lacked formal models for incentivizing diversity. The approach not only establishes theoretical guarantees but also introduces a game-theoretic framework for aligning incentives in decentralized settings, making it directly relevant to blockchain governance and infrastructure. Beyond improving resilience against censorship, natural disasters, or geopolitical disruptions, the ability to integrate location-aware guarantees into cryptographic protocols opens entirely new application domains for Cardano, including location-based key generation, enhanced staking mechanisms, and smart contract functionality tied to geographic constraints.

OO-1V: Ouroboros Peras – Vision

Start date: January 2025
Duration (months): 48

(i) Stream overview

Ouroboros Peras is a proof-of-stake consensus protocol that enhances the Cardano mainchain by reducing settlement times while preserving the robustness of the Ouroboros family. Whereas longest-chain protocols typically require many blocks to reach high-assurance finality, slowing adoption and complicating Layer 2 or cross-chain use, Peras enables fast settlement under normal conditions and falls back to current speeds during network instability or adversarial attacks.

A first version of the protocol has been specified, and served as a basis for an innovation stream evaluating its benefits as well as performance, showing promising estimated gains in settlement speed. In parallel, the research effort is focusing on providing formal security guarantees for the protocol, establishing its safety, liveness, self-healing properties, as well as a quantitative understanding of its settlement guarantees.

Further refinements will focus on improving efficiency of the protocol's cryptographic building blocks, improving its resilience in more adverse conditions, and addressing challenges such as avoiding cooldown phases, which remain non-trivial for longest-chain protocols. The ultimate goal is to deliver a production-ready protocol that strengthens the user experience on Cardano, supports advanced Layer 2 functionality, and enables more efficient bridging to other blockchains.

(ii) Context and objectives

While Cardano's current Ouroboros implementation is highly resilient and capable of operating under dynamic participation and adversarial stress, its gradual settlement model requires long confirmation times to achieve high assurance. This limits usability for end users and introduces friction for Layer 2 solutions and interoperability use cases that depend on fast, predictable finality. Our research aims to:

1. **Optimize settlement times** under typical conditions without weakening robustness under stress scenarios.
2. **Refine protocol parameters** and improve resilience to avoid cooldown phases, ensuring stability across varying network conditions.
3. **Produce formally verified research papers** that establish security proofs while balancing complexity and efficiency.
4. **Enable broader applications** of fast settlement, from on-chain user experience to Layer 2 deployments and secure cross-chain bridging.

By achieving these objectives, Peras will not only improve the day-to-day usability of the Cardano mainchain but also unlock new growth opportunities across the ecosystem, reinforcing Cardano's position as a secure, scalable, and interoperable blockchain platform.

(iii) Work performed and main achievements

The core of our work is the development of a scientific paper, *Adaptively Secure Fast Settlement with Dynamic Participation and Self-Healing, Christian Badertscher, Sandro Coretti-Drayton, Peter Gaži, Aggelos Kiayias, Alexander Russell,* which formally analyzes the Peras protocol—proving its safety, liveness, and self-healing properties. A manuscript is expected by the end of 2025. The protocol introduces periodic voting rounds, where participants vote on a block to settle; if enough votes are cast, a certificate is formed, boosting that block's weight. This shifts chain selection from "longest" to "heaviest," enabling significantly faster settlement under honest conditions. In cases of disagreement, a cooldown and restart mechanism ensures recovery and continuity.

The current version targets Cardano, using an "L slots deep" heuristic to select blocks for voting—offering optimistic performance as long as honest parties converge. Future work may replace this heuristic with more robust pre-agreement mechanisms. In parallel, we are optimizing committee sizes, reducing on-chain certificate storage, and improving communication efficiency in collaboration with the Leios and Mithril teams. Engineering efforts have also produced two technical reports: one demonstrating the feasibility of integrating Peras into Praos with minimal disruption, and another outlining interim mitigations that evolved into the Anti-Grinding workstream, now accompanied by a CPS, CPD, and CIP under submission.

Deliverables

#	Deliverable	Status
D1	Paper that formally analyzes the Ouroboros Peras protocol: *Adaptively Secure Fast Settlement with Dynamic Participation and Self-Healing*	In progress - expected '25
D2	Paper on modular improvements of the basic protocol (joint with Leios, Mithril)	In progress

(iv) Results beyond state of the art

A promising synergy has emerged among the Peras, Leios, and Mithril teams, centered on the need for a more compact and communication-efficient method for a majority of stakeholders to certify a given statement. The underlying primitive, *Stake-based Threshold Multisignatures*, was introduced in the original *Mithril: Stake-based Threshold Multisignatures*, *Pyrros Chaidos, Aggelos Kiayias, CANS '24*. However, it now appears feasible that its efficiency could be significantly improved using techniques from other ongoing research streams. Such improvements would benefit all three projects. In the case of Peras, they could be applied directly to the certificates that boost block weight, reducing both the communication overhead required to generate them (fewer parties need to vote and propagate their vote) and the certificate size itself (which is crucial when occasionally certificates are stored on-chain).

OO-2 Ouroboros Leios

Start date: January 2022
Duration (months): 42

(i) Stream overview

While Cardano's current consensus is robust, it is constrained by block size and timing, leaving most network resources such as bandwidth, CPU, and memory underused. Leios introduces a vertically scalable design, where throughput grows in proportion to node capacity—doubling resources roughly doubles performance.

It will be rigorously specified, proven secure in Cardano's security model, and supported by an incentive mechanism to ensure correct and efficient participation. By unlocking orders-of-magnitude higher throughput, Leios will enable the Cardano mainchain to support the growing demands of service chains, partnerchains, and high-volume smart contracts, laying the foundation for long-term scalability.

(ii) Context and objectives

Cardano's Praos protocol is reaching its throughput ceiling, with SPOs currently using less than 15% of their CPU resources. At the same time, growing adoption of smart contracts, DeFi, and sidechains requires significantly higher capacity. While existing research into throughput-optimal protocols shows promise, most rely on overly simplified models or apply only in permissioned settings. The main objectives are to:

1. **Design and analyze a scalable protocol** that achieves throughput proportional to network capacity.
2. **Address concurrency challenges** in high-rate settings.
3. **Ensure compatibility with Ouroboros Peras** to combine scalability with fast settlement.
4. **Introduce congestion control mechanisms** for stability under heavy load.
5. **Develop an incentive model** that rewards efficient participation.

Leios seeks to move beyond these limitations by delivering a provably secure, resource-efficient protocol tailored for permissionless networks. By meeting these goals, Leios will allow Cardano to scale far beyond current levels, strengthening its ability to support a diverse and growing ecosystem.

(iii) Work performed and main achievements

The Leios research has made significant progress in 2025 toward its goal of scaling throughput in permissionless settings. The first major milestone—design and analysis of a vertically scalable blockchain protocol—has been completed, with results published as *High-Throughput Permissionless Blockchain Consensus under Realistic Network Assumptions*, *Sandro Coretti-Drayton, Matthias Fitzi, Aggelos Kiayias, Giorgos Panagiotakos, Alexander Russell,* at Crypto '25. This work demonstrates that throughput can scale in line with available resources, addressing the core limitation of Praos.

Building on this foundation, the team has begun investigating concurrency issues that arise when multiple blocks are produced in parallel. Early work focuses on strategies to minimize or prevent the inclusion of conflicting transactions in concurrently generated blocks, with a dedicated paper in preparation. Additional tasks on compatibility with Ouroboros Peras and congestion control mechanisms are scheduled to follow later in the workplan.

Deliverables

#	Description	Status
D1	Paper on how how to scale data throughput: *High-Throughput Permissionless Blockchain Consensus under Realistic Network Assumptions*	Published at Crypto '25
D2	Paper on concurrency issues such as minimizing the inclusion of conflicting transactions in produced blocks	In progress

(iv) Results beyond state of the art

Leios sets a new benchmark for blockchain scalability by delivering the first provably secure, high-throughput protocol for permissionless networks under realistic assumptions. Prior works claiming throughput optimality have relied on idealized models or permissioned settings, leaving open the challenge of achieving scalability with adaptive security in open environments. The Leios paper presented at Crypto 2025 addressed this gap by introducing a realistic gossip-based network model that captures real-world challenges such as message bursts and proof-of-stake equivocations, which existing models largely ignore.

Building on this model, Leios transforms any base protocol into one achieving throughput close to a (1–δ) fraction of network capacity, while preserving constant expected settlement time when available. This positions Leios well ahead of alternatives like Narwhal/Tusk, which lack proofs for fully decentralized settings. Going forward, research will expand to concurrency handling, network-layer redesign, and congestion control—areas that remain unsolved in the literature. Together, these advances could redefine efficiency and fairness in proof-of-stake consensus and further strengthen Cardano's leadership in secure, scalable blockchain protocols.

OO-3 Fair transaction processing

Start date: Jul 24
Duration (months): 42

(i) Stream overview

The Fair Transaction Processing workstream aims to address a long-standing limitation of blockchain systems: the lack of guarantees around transaction ordering. Like Bitcoin and Ethereum, Cardano's current protocol allows block producers to order transactions arbitrarily, enabling harmful practices such as front-running and Maximal Extractable Value (MEV). These practices can distort markets, reduce trust, and harm end users.

This project is developing protocol-level solutions that integrate fairness into the Ouroboros consensus family. By combining cryptographic techniques with consensus modifications, the goal is to provide equitable transaction ordering by default, without compromising decentralization or performance. If successful, Cardano would become the first major blockchain to offer provable fairness guarantees at the protocol layer—an innovation that directly benefits end users and strengthens Cardano's competitive position.

(ii) Context and objectives

MEV and unfair ordering have become central issues in blockchain design. Existing mitigations, such as proposer-builder separation, can reduce MEV but introduce new risks of centralization, undermining decentralization goals. Cardano's research agenda instead seeks a first-principles solution: embedding fairness directly into consensus, where it cannot be bypassed by external actors. The team has already made progress in this area, particularly in formalizing fairness definitions and proposing new constructions. The specific objectives of the stream are to:

1. **Design and analyze a consensus extension** for Ouroboros that enforces fair transaction processing under formal definitions of fairness.
2. **Ensure compatibility with Peras and Leios**, so that fairness is preserved in the context of fast settlement and high throughput.

These works position Cardano at the forefront of the global research effort to solve this problem. If achieved, this work would represent a breakthrough for blockchain usability and trust, making fairness a unique feature of Cardano's consensus protocol and a strong differentiator in an increasingly competitive ecosystem.

(iii) Work performed and main achievements

At Crypto '24, the team presented *Universal Composable Transaction Serialization with Order Fairness*, *Michele Ciampi, Aggelos Kiayias, Yu Shen*, establishing impossibility results and introducing a practical enclave-based protocol that preserves input causality while achieving optimal sender fairness. At Eurocrypt '24, *Ordering Transactions with Bounded Unfairness: Definitions, Complexity and Constructions*, *Aggelos Kiayias, Nikos Leonardos, Yu Shen* defined the notion of bounded unfairness, proved its inherent complexity, and introduced Taxis, a protocol that balances near-optimal fairness with liveness and performance trade-offs.

In 2025, the Fair Transaction Processing stream achieved further progress with *Universally Composable Transaction Order Fairness: Refined Definitions and Adaptive Security*, *Michele Ciampi, Aggelos Kiayias, Yu Shen*, accepted at Asiacrypt '25, which refines the UC framework by unifying receiver-order fairness and input causality, incorporating transaction fees, and introducing a novel YOSO-style encryption scheme that achieves fairness with decryption complexity independent of the number of transactions—the first such result under standard cryptographic assumptions. Collectively, these works provide a rigorous foundation and practical pathways for ensuring fairness in blockchain transaction processing.

Deliverables

#	Description	Status
D1	Paper on fair transaction processing networking layer support: *Universally Composable Transaction Order Fairness: Refined Definitions and Adaptive Security*	Accepted Asiacrypt '25
D2	Paper on fair transaction processing overlay for Ouroboros: design and analysis	In progress

(iv) Results beyond state of the art

This work establishes the first formal and practical framework for fair transaction ordering in blockchains—a capability no public network currently offers. The research unifies receiver-order fairness and input causality, incorporates transaction fees, and introduces a novel YOSO-style encryption scheme that enables fair ordering with decryption complexity independent of transaction volume. This breakthrough positions Cardano to lead in provably fair transaction processing, offering a unique and verifiable advantage in blockchain infrastructure.

OO-5 Multi-resource consensus - Minotaur

Start date: Jan 25
Duration (months): 24

(i) Stream overview

Minotaur explores the design of consensus protocols that draw on multiple resources—such as proof of stake (PoS), proof of work (PoW), and proof of restake (PoRS)—to improve security, resilience, and inclusivity. By combining diverse resources, Minotaur ensures that consensus integrity does not depend on any single factor: even if one resource is weak or under attack, the system can continue to function securely.

This approach is particularly valuable for bootstrapping new blockchains, such as PartnerChains, where liquidity or stake distribution may be low at launch. By leveraging existing assets—for example, restaked PoS from the Cardano mainchain—emerging networks can rely on a stable base of participants to provide security from day one. Minotaur represents a flexible, hybrid consensus model designed to strengthen the robustness of decentralized systems and expand the utility of Cardano's ecosystem.

(ii) Context and objectives

Most consensus protocols rely on a single resource—PoW in Bitcoin or PoS in Cardano—making them vulnerable if that resource is scarce or concentrated, especially in early-stage networks. The original Minotaur paper showed how PoW and PoS could be securely combined, laying the groundwork for broader multi-resource consensus. This stream extends the approach to BFT consensus and PartnerChains, formally defining how diverse resources—PoS, PoW, PoRS, or useful work—can be measured, weighted, and combined into a unified metric for block production. Key objectives are to:

1. **Design a provably secure multi-resource consensus protocol** for sidechains, ensuring resilience even if one resource is compromised.
2. **Enable bootstrapping of PartnerChains** by combining Cardano stake and SPO participation with local tokens for immediate security.
3. **Demonstrate performance improvements** compared to single-resource BFT consensus and validate scalability for practical deployment.

By achieving these goals, Minotaur will provide the first formal framework for multi-resource consensus, making it easier to launch secure, inclusive PartnerChains and strengthening Cardano's position as a platform for decentralized innovation.

(iii) Work performed and main achievements

We aim towards a technical report or paper about multi-resource consensus for BFT (i.e., committee-based) sidechains, involving re-staking from different blockchain, together with security proofs. A first protocol version was sketched with a focus on integrating Ethereum re-staking as a resource (alongside with other resources), and applying randomized committee selection by using pseudo-randomness generated on the sidechain.

A second protocol version is currently being developed with a focus on integrating Bitcoin re-staking along the lines of Babylon Chain (together with other resources), and composing the sidechain committees by deterministic selection from the top "contributors" from each involved resource. Improvements over the guarantees given by Babylon Chain are being explored for Babylon Chain itself (i.e., in the single resource case), and for the extension to multiple resources, e.g., examining under which conditions liveness can be regained after it has been lost.

Deliverables

#	Description	Status
D1	Paper describing a provably secure implementation of multi-resource BFT consensus or its components	In progress - expected '25

(iv) Results beyond state of the art

We plan to provide the first BFT consensus protocol driven by multiple resources and integrating with major blockchains such as Bitcoin and Ethereum. Furthermore, we hope to improve over the security guarantees for Bitcoin re-staking as provided by Babylon Chain promising variants that offer better safety of actively validated services driven by stake in Bitcoin.

OO-6 Proofs of useful work

Start date: Jan 21
Duration (months): 60

(i) Stream overview

Proofs of Useful Work (PoUW) aim to transform blockchain consensus by replacing wasteful Proof-of-Work with real-world computational tasks that deliver value beyond network security. Building on *Ofelimos: Combinatorial Optimization via Proof-of-Useful-Work*, *Fitzi et al.,* (Crypto '22), Cardano's first provably secure PoUW protocol, this research explores extending the model to optimization problems, zk-SNARK generation, and machine learning.

By securing Cardano through useful computation, PoUW not only improves sustainability but also creates an on-ramp for new participants whose contributions generate tangible scientific and industrial benefits. PoUW can also reinforce multi-resource consensus (e.g., Minotaur) and enhance blockchain efficiency, such as enabling faster light-client validation via SNARKs.

(ii) Context and objectives

Traditional PoW wastes energy on arbitrary puzzles, while existing PoUW proposals often lack rigorous security guarantees. Ofelimos demonstrated that optimization problems can be embedded into consensus with formal security proofs, but challenges remain in generalizing the approach, ensuring robustness across diverse tasks, and aligning incentives. This research therefore sets out to:

1. **Advance optimization-based PoUW** beyond Ofelimos, ensuring both blockchain security and computational performance.
2. **Extend PoUW to new domains** such as SNARK generation and machine learning.
3. **Develop incentive and economic models** that align miners' rewards with useful outcomes.
4. **Explore privacy-preserving features** for PoUW blockchains.

By addressing these objectives, PoUW can make Cardano the first major blockchain ecosystem to integrate provably secure, useful computation into its consensus—unlocking both stronger security and real-world impact.

(iii) Work performed and main achievements

Recent research has built on the foundation laid by Ofelimos. A follow-up paper, *Proof-of-Useful-Work is Practical: Consensus via a Competitive Optimization Engine*, *Matthias Fitzi, Aggelos Kiayias, Laurent Michel, Giorgos Panagiotakos, Alexander Russell* (under submission), advances this work by demonstrating how PoUW can be implemented efficiently in real-world consensus settings. Together, these contributions establish a rigorous basis for embedding optimization tasks into blockchain security.

On the machine learning side, exploratory work tested the feasibility of Proof-of-Deep-Learning (PoDL) as a follow up to *Provably Secure Blockchain Protocols from Distributed Proof-of-Deep-Learning*, *Xiangyu Su, Mario Larangeira, and Keisuke Tanaka, 2023*. A transformer-based predictor with a stabilization algorithm was developed as part of a university collaboration, extending the existing 2023 prototype, to detect when training ceases to be useful, addressing challenges like loss saturation and dataset availability. Validated on tasks such as MNIST and CIFAR-10, this approach was released in an *open-source repository*, offering a flexible framework for PoUW in deep learning. This marks an initial step toward extending PoUW beyond optimization into broader computational domains.

Deliverables

#	Description	Status
D2	Paper on applicability of PoUWs to other domains: *Proof-of-Useful-Work is Practical: Consensus via a Competitive Optimization Engine*	Under submission '25

(iv) Results beyond state of the art

PoUW advances blockchain consensus by securing networks through meaningful computation rather than wasteful hashing. Building on *Ofelimos* and subsequent work, it is among the first frameworks to offer formal security guarantees for optimization-based PoUW, with extensions into machine learning via Proof-of-Deep-Learning (PoDL).

This work addresses open challenges such as detecting useful training phases, dataset availability, and incentive design—pushing PoUW from theory toward practical deployment and expanding its applicability across optimization, SNARK generation, and ML tasks.

OO-7 Congestion control

Start date: Jan 22
Duration (months): 60

(i) Stream overview

This research stream rethinks blockchain fee mechanisms to deliver fairer and more predictable congestion control on Cardano. Existing flat or dynamic fees either lack flexibility or create uncertainty, limiting support for diverse use cases such as DeFi, payments, and high-priority transactions.

The research explores models where fees reflect both resource usage and urgency, allowing users to choose trade-offs between cost and speed. Building on *Tiered Pricing*, the aim is to design a modular fee system that ensures efficient resource allocation, reduces spam vulnerability, and improves user experience under varying network conditions.

(ii) Context and objectives

Congestion control is critical to ecosystem growth: predictable fees are essential for businesses planning transactions in advance, while flexibility supports a wide range of applications with different latency and cost requirements. Current blockchain fee designs, though studied extensively, have yet to produce a fundamentally superior model. This research sets out to:

1. **Expand congestion control** to handle application diversity and user urgency.
2. **Improve predictability** by offering service guarantees on cost and delay.
3. **Address resource heterogeneity** by pricing bandwidth, storage, and other resources separately.
4. **Develop parameter update processes** that adapt quickly to demand while maintaining stability.

The outcome will be a fee architecture that enables greater participation, strengthens business adoption, and optimizes Cardano's ability to handle varied workloads at scale.

(iii) Work performed and main achievements

This research has already produced two peer-reviewed publications: *Tiered Mechanisms for Blockchain Transaction Fees*, *Kiayias et al.* (MARBLE '24) and *Blockchain Space Tokenization*, *Kiayias et al.* (AFT '24). Current work is advancing along three active research tracks: (1) designing predictable service mechanisms via blockchain space futures (building on Cardano's fixed-fee model), (2) developing methods to handle multi-resource congestion, and (3) unifying congestion control with tokenomics design.

One-dimensional vs. Multi-dimensional Pricing in Blockchain Protocols, Aggelos Kiayias, Elias Koutsoupias, Giorgos Panagiotakos, Kyriaki Zioga, is under submission, which analyzes the trade-offs between simple, fast-converging fee models and more efficient but complex multi-dimensional pricing. In parallel, a master thesis *Service Predictability in Blockchain, Georgios Tsironis (and in collaboration with Aristides Pagourtzis, Aggelos Kiayias, Elias Koutsoupias, Giorgos Panagiotakos)* has been completed, proposing a derivative-based protocol to guarantee predictable transaction servicing. At least one additional paper is expected to emerge from this line of research, specifically addressing resource heterogeneity.

Deliverables

#	Description	Status
D2	(b) Paper on auctions and predictable delays (c) Paper on blockchain space derivatives: *Service Predictability in Blockchain*	(b) In progress (c) Thesis published ‘25
D3	(a) Paper on multi-resource congestion: *One-dimensional vs. Multi-dimensional Pricing in Blockchain Protocols* (b) Paper on semi-permanent resources	(a) Under submission ‘25 (b) Not yet started

(iv) Results beyond state of the art

This workstream pushes blockchain fee and congestion models beyond existing designs, which remain largely flat or one-dimensional. The tiered pricing framework introduces a flexible extension to EIP-1559, enabling users to choose between cost and latency trade-offs, something not supported in current systems.

The forthcoming work on multi-dimensional pricing is the first to rigorously analyze the welfare and stability trade-offs of pricing heterogeneous resources, offering theoretical foundations for resource-aware fee markets. The blockchain space futures research introduces an entirely novel approach, using financial derivatives to guarantee block space predictability, unachievable in dynamic fee systems like Ethereum.

Together, these results set a new standard for congestion control by combining economic efficiency, predictable service, and support for heterogeneous resources, paving the way for more reliable and diverse blockchain applications.

TO-1 Tokenomics design

Start date: Jan 25
Duration (months): 48

(i) Stream overview

The Tokenomics Design stream investigates how to build a sustainable and resilient economic foundation for Cardano through first-principles research. By developing rigorous mathematical models, the project will capture the interplay between core system components such as staking rewards, treasury mechanisms, reserves, and transaction fees, as well as broader factors like adoption and external shocks. These models aim to reflect the heterogeneous and decentralized nature of the Cardano ecosystem, enabling precise evaluation of long-term stability and equilibrium.

The ultimate goal is to provide evidence-based guidance for Cardano's macroeconomic policy design, helping ensure that ada retains strong value, supports broad participation, and continues to attract users and developers. In doing so, the research contributes not only to the network's security and robustness but also to its capacity for sustainable growth and innovation.

(ii) Context and objectives

Most existing tokenomics frameworks focus narrowly on either user adoption dynamics or validator incentives, leaving significant gaps in modeling whole-ecosystem effects. Cardano's design—with features like reward sharing, treasury funding, and reserves—requires a more comprehensive approach that integrates multiple perspectives into a unified model. This research addresses that gap, examining how diverse system parameters collectively shape price evolution, security, and community-led governance. The key objectives are to

1. **Identify the parameters** most critical to token price and long-term equilibrium
2. **Explore equilibrium dynamics** under stress conditions such as liquidity shocks or adoption shifts
3. **Propose optimal settings** and governance mechanisms that balance resilience with growth.

By delivering these insights, the project equips the Cardano community with tools to make informed, forward-looking decisions that sustain security while unlocking greater economic and innovation potential.

(iii) Work performed and main achievements

Over the past year, research in this stream has advanced the theoretical and practical foundations of blockchain tokenomics with a focus on Cardano's decentralized and heterogeneous ecosystem. A major achievement was the development of new models for assessing long-term system equilibria, accounting for reserve depletion, parameter trade-offs, and the interplay between adoption, participation, and economic security. This work builds on *Single-token vs Two-token Blockchain Tokenomics*, *Aggelos Kiayias, Philip Lazos, Paolo Penna*, accepted at AFT '25, which provides a formal framework for analyzing token economies. A follow-up study, now in preparation, extends this model to Cardano-specific scenarios such as reserve policy impacts on token price evolution and system resilience under external shocks.

Complementing the theoretical work, an innovative experimental game-theory approach was introduced to capture real-world crypto-user behaviors. Large-scale surveys across four countries with over 11,000 participants yielded first-of-its-kind data on adoption, risk preferences, and time horizons, distinguishing crypto holders from stock investors. These findings will be incorporated into refined mathematical models to more accurately reflect real user dynamics in tokenomics design. In parallel, the team developed a novel framework *Algorithmic Monetary Policies for Blockchain Participation Games*, *Diodato Ferraioli, Manvir Schneider, Paolo Penna*, for analyzing policies and their induced trade-offs, presented at the 2025 Archimedes Workshop.

Deliverables

#	Description	Status
D1	Paper on long-term equilibria: (i) *Single-token vs Two-token Blockchain Tokenomics*	(i) Accepted AFT '25
D2	Paper analyzing monetary policies: *Algorithmic Monetary Policies for Blockchain Participation Games*	Under submission '25
D3	Paper on Cardano's key parameters	In progress
D4	Paper on the ecosystem as a whole, including side-chains and multi-token platforms	In progress

(iv) Results beyond state of the art

This research advances tokenomics by moving beyond traditional models that assume infinite horizons or negligible individual influence. Our new frameworks explicitly incorporate finite launch phases, heterogeneous participants, and ecosystem-wide effects—factors critical for real-world blockchain sustainability. This enables more accurate evaluation of policies such as reserve depletion schedules and adaptive rewards, offering practical guidance for Cardano's economic design.

A major innovation is the integration of *experimental game theory* into tokenomics. Large-scale surveys have grounded models in observed crypto-user behavior, revealing differences from traditional investors and highlighting new levers for adoption and participation incentives. Complementary work on airdrop mechanisms establishes launch-phase tokenomics as a distinct research field, directly informing partnerchain and DApp rollouts. Collectively, these results position Cardano at the forefront of principled, empirically informed tokenomics—treating it as a genuine "monetary policy for decentralized systems" that shapes stability, adoption, and trust.

TO-2 Rewards sharing and transaction fees

Start date: Jun 23
Duration (months): 48

(i) Stream overview

Effective reward distribution is central to sustaining a secure and decentralized blockchain. In Cardano, this depends on carefully designed schemes that share block production rewards and transaction fees between stake pool operators (SPOs) and delegators. However, designing such mechanisms is challenging: they must balance fairness and efficiency, account for asymmetries among participants, and operate under anonymity constraints. Poorly tuned mechanisms risk centralization or reduced user satisfaction, while robust, fair, and transparent schemes strengthen both security and ecosystem adoption.

This workstream develops improved reward sharing and fee models using advanced game-theoretic analysis. By refining the incentive structure, the project aims to ensure Cardano remains attractive for both service providers and end users—delivering equitable returns for participants, predictable and fair fees, and enhanced decentralization. The research will not only inform upgrades to Cardano's existing system but also guide reward models for governance, Project Catalyst, and protocols like Mithril.

(ii) Context and objectives

While game theory has been widely applied to analyze incentives in economic systems, only limited work has been tailored to the unique requirements of blockchain environments. Cardano already benefits from pioneering reward-sharing research, but open questions remain around fairness, variance, long-term decentralization, and the design of incentive mechanisms across different system layers. Addressing these gaps is critical to sustaining healthy participation, ensuring diverse stake distribution, and keeping Cardano competitive as a platform for decentralized services. The objectives of this research are therefore threefold:

1. **Deepen the theoretical understanding** of blockchain-specific reward dynamics.
2. **Propose and evaluate new mechanisms** that enhance fairness, decentralization, and predictability.
3. **Extend these insights to other Cardano components**, including DRep rewards under Voltaire and reward structures in Catalyst and Mithril.

Ultimately, the outcomes will provide both rigorous academic insights and practical design improvements for Cardano's incentive layer.

(iii) Work performed and main achievements

During the first half of 2025, our work centered on exploring alternative reward-sharing mechanisms and broadening their applications within the Cardano ecosystem. A major outcome was the development of a Shapley-value–based scheme for on-chain pooling, which offers a principled approach in allocating rewards as an alternative solution to proportional sharing. This research culminated in the paper *Pool Formation in Oceanic Games: Shapley Value and Proportional Sharing*, *Aggelos Kiayias, Elias Koutsoupias, Evangelos Markakis, Panagiotis Tsamopoulos*, accepted for presentation at AFT '25.

In parallel, some initial progress was made in comparing the trade-offs between on-chain and off-chain pooling models, with the aim of creating more accurate mathematical frameworks for evaluating their benefits and drawbacks. We also initiated the design of incentive mechanisms for other Cardano applications, such as Mithril certificate signing, where robust participation and fair distribution of rewards are critical. Dissemination of results has included a keynote at the workshop ALGA '25 and participation at the conference FC '25, ensuring active engagement with the wider research community.

Deliverables

#	Description	Status
D1	Paper on the study of alternative reward schemes: *Pool Formation in Oceanic Games: Shapley Value and Proportional Sharing*	Accepted AFT '25
D2	Paper on on-chain vs off-chain pooling	In progress

(iv) Results beyond state of the art

This research pushes beyond the state of the art by introducing the Shapley value, a well-established concept in cooperative game theory, into blockchain reward-sharing. While proportional reward allocation dominates in current systems, the Shapley-based approach captures each participant's true marginal contribution, offering guarantees of fairness and long-term stability. This marks a significant step in bridging theoretical economics with practical blockchain incentive design.

By extending the analysis beyond staking pools to include off-chain pooling models and protocols such as Mithril, the work establishes new ground for incentive engineering in heterogeneous blockchain environments. The result is not only a more sophisticated understanding of how rewards can drive fairness and efficiency in Cardano, but also a reusable framework for incentive design across diverse decentralized systems.

D4-1 Next-level governance protocols

Start date: Jan 24
Duration (months): 48

(i) Stream overview

This workstream advances Cardano's governance beyond CIP-1694 by developing scalable, decentralized voting protocols that remain reliable under high network activity. Existing systems either depend on centralized tallying or are too costly to run fully on-chain. To overcome these limits, the project explores lightweight, privacy-preserving, and verifiable protocols that balance on-chain transparency with off-chain and Layer 2 efficiency.

By designing modular governance mechanisms—ranging from stake-weighted voting to representative and liquid democracy models—this research aims to equip Cardano with future-proof governance tools that support broad participation, ensure strong security guarantees, and evolve with the ecosystem's growth.

(ii) Context and objectives

Governance is central to blockchain sustainability, yet most e-voting protocols fail to meet the needs of large-scale, permissionless systems. They either rely on trusted authorities for tallying or are too costly to run efficiently on decentralized networks. Cardano must overcome these limitations to enable inclusive participation while preserving decentralization and constitutional alignment. Prior work, including treasury-based voting and stake-based multisignature schemes, offers useful foundations but does not fully resolve the challenge of combining scalability, security, and cost-efficiency. The objectives of this research are fourfold:

1. **Systematically capture the requirements** for governance protocols across Cardano use cases.
2. **Design protocols** that ensure privacy, verifiability, and censorship resistance
3. **Evaluate trade-offs** in performance, cost, and scalability across direct and delegated voting models.
4. **Provide formal guarantees** alongside practical, low-footprint implementations.

Success would equip Cardano with a next-generation governance toolkit—capable of supporting Catalyst, foundation elections, and broader governance efforts—with models that can be extended to other ecosystems in the future.

(iii) Work performed and main achievements

In the first half of 2025, the research advanced the formal foundations of governance within Cardano's ecosystem by finalizing a Universally Composable (UC) security framework tailored to layer-2 governance protocols. This framework builds on three previously introduced ideal functionalities—Distributed Homomorphic Encryption with Aggregation (F-DHE), Ballot Aggregation with Validity Proofs (F-BB), and Secure Voting with Dispute Resolution (F-L2Gov)—to provide a robust basis for analyzing governance mechanisms in decentralized systems.

Leveraging this framework, the team designed a new governance protocol that achieves its security guarantees with minimal layer-2 footprint and minimal assumptions, making it well-suited to Cardano's heterogeneous and resource-conscious environment. The protocol and accompanying model were consolidated into *Beyond Blockchain Ballots: UC-Secure Layer-2 Voting and Governance*, *Raghav Bhaskar, Pooya Farshim, Matthias Fitzi, Aggelos Kiayias*, which is under submission. This milestone represents the key planned deliverable for 2025, combining the intended contributions of both protocol design and formalization.

Deliverables

#	Description	Status
D1	Paper on the trade-offs of security and cost when using various layer 2 primitives to lower the overall governance cost	In progress
D2	Paper that brings out the benefits of using a blockchain for a governance protocol: *Beyond Blockchain Ballots: UC-Secure Layer-2 Voting and Governance* *	Under submission '25
D3	The security framework developed to study the security of the various protocols may itself be worth publishing	In progress

* This paper is a joint deliverable of D2 and D3

(iv) Results beyond state of the art

This work advances the state of the art in decentralized governance by delivering the first formal UC-based security model tailored specifically to layer-2 governance protocols. Unlike prior approaches that rely on ad hoc security arguments or heavyweight assumptions, this model provides composable guarantees and supports modular reasoning about governance in complex ecosystems such as Cardano.

The newly proposed protocol is equally notable: it achieves strong security properties while minimizing costs in terms of computation, communication, and blockchain footprint. This makes it practical for real-world deployment while maintaining rigorous cryptographic guarantees. By combining foundational formalization with deployable protocol design, the research sets a new benchmark for secure, efficient, and scalable governance mechanisms in blockchain systems.

D4-2 Governance incentives

Start date: Jan 23
Duration (months): 48

(i) Stream overview

Designing effective governance incentives is essential for balancing fast, impactful decision-making with decentralization, fairness, and security. Current models that rely mainly on token ownership risk centralization and fail to capture the broader diversity of participants.

This workstream focuses on designing robust incentive mechanisms, starting with Delegated Representatives (DReps) and extending to other actors in the ecosystem. By aligning incentives with honest participation and long-term ecosystem health, this research aims to create governance structures that are more inclusive, transparent, and resistant to manipulation.

(ii) Context and objectives

The research will investigate alternatives to the "1 token, 1 vote" model, evaluating approaches that achieve better proportional representation and better reflect voter preferences, mitigating centralization risks. It will also study treasury fund allocation mechanisms that balance budget efficiency with fairness, ensuring that decision-making is not dominated by simple majorities.

Practical application will target Voltaire and Project Catalyst, producing more expressive voting rules and improved incentive structures. The objectives are to deepen understanding of governance issues, adapt cutting-edge participatory budgeting methods, and design reward schemes that ensure DReps and other participants are motivated to act truthfully and responsibly.

(iii) Work performed and main achievements

In the first half of 2025, this stream advanced research on incentive structures and voting mechanisms central to Cardano's governance. A key milestone was the publication of *Reward Schemes and Committee Sizes in Proof of Stake Governance*, *Georgios Birmpas, Philip Lazos, Evangelos Markakis, Paolo Penna*, (FC '25), which develops a formal model of delegated representatives (DReps) where effort, delegation, and rewards jointly shape governance outcomes. This work provides principled insights into how reward mechanisms and committee composition affect both decision accuracy and robustness, offering concrete guidance for incentive design in Voltaire.

Parallel efforts explored novel approaches to participatory budgeting, including models where community donations supplement fixed budgets, with applicability to Catalyst. The research team has also engaged widely with the academic community, presenting results at Financial Cryptography (FC '25) and the Workshop on Algorithms, Learning, and Games (ALGA '25), ensuring Cardano's governance research is visible to both cryptography and economics communities.

Deliverables

#	Description	Status
D1	Paper on DRep incentives and reward schemes: *Reward Schemes and Committee Sizes in Proof of Stake Governance*	Published FC '25
D2	Paper on participatory budgeting	In progress - expected '25

(iv) Results beyond state of the art

This research pushes beyond existing literature by introducing a rigorous game-theoretic framework for analyzing governance incentives in delegated proof-of-stake systems. Previous work on blockchain governance has largely focused on descriptive or empirical models; in contrast, this work formalizes the trade-offs between DRep effort, committee composition, and reward allocation under constrained budgets. The resulting insights provide a principled basis for optimizing governance incentives, ensuring both accuracy in decision-making and robustness against misaligned incentives.

The participatory budgeting research further extends the state of the art by modeling hybrid funding mechanisms where community donations complement treasury resources—a novel perspective in blockchain governance design. Together, these contributions position Cardano at the forefront of incentive-aligned, scientifically grounded governance, directly informing the rollout of Voltaire and the evolution of Catalyst.

IHT-1 Hydra tail

Start date: Jan 22
Duration (months): 36

(i) Stream overview

Hydra tail builds on Cardano's Hydra framework by introducing zk-rollups as a complementary scaling solution to Hydra Heads. Instead of relying solely on state channels, Hydra Tail batches many off-chain transactions and posts succinct zero-knowledge proofs back to the main chain. This reduces on-chain load, preserves strong security guarantees, and enables efficient fund and contract mobility between layer 1 and layer 2.

The protocol is designed around a state-channel (SC) model, where channels can open, close, fail, or resolve disputes, ensuring predictable outcomes. A key feature is the use of registered transactions (regTx), which guarantee censorship resistance by allowing users to safely reclaim funds, even if operators attempt to block participation.

(ii) Context and objectives

The primary goal is to design a zk-rollup protocol from first principles and prove its security within the UC framework. This involves formally defining the ideal functionalities a secure zk-rollup should provide—such as verifiability, data availability, and censorship resistance—before proposing a concrete design tailored for UTxO-based ledgers.

The research will balance three critical objectives: minimizing mainchain footprint, ensuring strong security guarantees, and maintaining protocol simplicity. By doing so, Hydra Tail aims to deliver a scalable and provably secure zk-rollup that enhances Cardano's throughput and usability without compromising decentralization.

(iii) Work performed and main achievements

The stream has developed a rollup-style multi-party state channel protocol, with the design and high-level code (including smart contract models) largely complete. In this approach, an untrusted server manages the off-chain ledger while security is anchored on Cardano's main chain: the server confirms off-chain transactions to the sender and periodically posts the updated state on-chain with a proof of correctness.

The protocol is designed to be uncollateralized, provide execution awareness through receipts, and include protections against malicious servers. While every off-chain transaction must eventually appear on-chain—a known trade-off in EUTxO-style systems—the model remains practical. The next milestone is completing a research paper with a full security proof, planned by the end of 2025.

Deliverables

#	Description	Status
D1	Paper on security proofs	In progress - expected '25

(iv) Results beyond state of the art

ZK-Rollups are important components in many blockchains. Yet, surprisingly, no formal cryptographic security analysis has been published for any of those solutions. We aim at presenting a ZK-rollup construction and prove it secure in the universal-composability framework by Canetti. Note that the focus here is on proving security rather than coming up with fundamentally new constructions.

IHT-2 Inter-head

Start date: Jan 25
Duration (months): 60

(i) Stream overview

The Inter-head protocol extends Cardano's Layer 2 scaling solutions by connecting multiple Hydra heads into a network of interoperable state channels. This enables participants to transact off-chain across different heads, combining the efficiency of Hydra with the flexibility of multi-party channel networks.

While Hydra and Inter-head significantly improve scalability by reducing on-chain footprint, their adoption for complex applications depends on rigorous assurances of security. To this end, the project applies the Universal Composability (UC) framework to formally analyze Hydra and Inter-head, ensuring that they can be safely composed with other protocols and used as reliable middleware for large-scale decentralized applications.

(ii) Context and objectives

Layer 2 protocols are widely regarded as a promising approach to blockchain scalability, but most existing designs lack comprehensive formal security treatment. The UC framework provides the state-of-the-art methodology for proving security of cryptographic protocols in composable environments, making it the natural choice for analyzing Hydra and Interhead. The objectives of this project are to:

1. **Formalize the ideal functionality** of Hydra and Interhead within the UC framework.
2. **Prove their security** under rigorous adversarial models, ensuring properties like liveness and safety hold in all contexts.
3. **Establish a foundation for building** secure, complex applications on top of Hydra-based networks by enabling safe composition with other Layer 2 and Layer 1 protocols.

This work will lay the groundwork for Hydra and Interhead to evolve into provably secure infrastructure, capable of supporting high-throughput applications while maintaining the robustness required for Cardano's long-term growth.

(iii) Work performed and main achievements

In 2025, significant progress was made toward establishing a composable security foundation for Hydra and Inter-head. The team developed a full Universally Composable (UC) security definition for multiparty isomorphic state channels, together with a functionality that captures their core properties. Building on this framework, they proved the security of Hydra (and related protocols) within the UC setting, thereby providing the first rigorous composable guarantees for the protocol. This work has already resulted in a paper *Universally Composable Treatment of Multi-Party Isomorphic State Channels*, *Maxim Jourenko, Xiangyu Su, Adam Blatchley Hansen, Mario Larangeira,* which is under submission to an international conference and ePrint.

On the Interhead side, preparatory work has begun to extend the UC framework to cover multiparty state channels that can be merged—an essential requirement for scaling Hydra through channel composition. The forthcoming deliverables include a formal UC definition for Interhead, a proof of its security under this model, and subsequent submissions to a top-tier conference planned for 2026.

Deliverables

#	Description	Status
D1	Paper on formalization of Hydra in the UC framework: *Universally Composable Treatment of Multi-Party Isomorphic State Channels*	Under submission '25
D2	Paper on formalization of the Interhead Protocol in the UC framework	In progress

(iv) Results beyond state of the art

This research marks the first universally composable treatment of Hydra, elevating it from an efficient layer-2 protocol to one with formally proven, composable security guarantees. Prior analyses of state channels often relied on bespoke or non-composable arguments, leaving open questions about security in complex multi-protocol environments. By introducing a UC framework for isomorphic multiparty state channels, this work enables modular reasoning and ensures Hydra's security even when composed with other protocols.

The forthcoming Inter-Head extension pushes the frontier further by tackling the unique challenges of mergeable multiparty channels, an area with no established formal treatment in the literature. Establishing a UC-secure model for Interhead will set a new benchmark for the scalability and composability of state channel protocols, positioning Cardano's layer-2 ecosystem at the forefront of both theoretical rigor and practical applicability.

IHT-3 Optimization tools

Start date: Jul 25
Duration (months): 60

(i) Stream overview

This research stream develops optimization tools to enhance the performance of Hydra-based Layer 2 networks. While Hydra provides scalable off-chain transaction processing, practical inefficiencies such as liquidity imbalances, routing bottlenecks, and channel synchronization issues can limit throughput and reliability.

By designing companion protocols that address these challenges, the project aims to ensure smoother operation, higher resource utilization, and greater resilience in real-world, high-demand scenarios. These tools will extend Hydra's scalability benefits, improve user experience, and prepare the protocol suite for broader adoption across diverse applications.

(ii) Context and objectives

Layer 2 protocols often face secondary challenges beyond their core design, particularly in managing liquidity flows and coordinating transactions across multi-party networks. Existing surveys of layer 2 tools highlight potential approaches, but many are outdated or not directly compatible with Hydra's architecture.

This project adopts a three-phase approach: first, surveying and comparing existing optimization tools in other protocols; second, evaluating their applicability to Hydra's model; and third, designing new protocols to fill the identified gaps. Objectives include developing mechanisms for fund rebalancing, efficient routing, and synchronization, while formally ensuring their compatibility and composability with Hydra and Interhead.

(iii) Work performed and main achievements

In the second half of 2025, a survey of the landscape of optimization tools for layer-2 protocols started with the aim of identifying mechanisms that could enhance Hydra's efficiency. The team systematically reviewed existing solutions for challenges such as fund rebalancing, message routing, and channel synchronization, comparing them to Hydra's specific requirements. This survey culminated in a comprehensive internal report that catalogues available approaches, highlights gaps in current tooling, and provides a structured mapping of which techniques may be adaptable to Hydra's architecture.

This foundational work not only established a clear picture of the state of the art but also created a roadmap for subsequent stages of the stream. It sets the stage for evaluating the compatibility of promising tools with Hydra in 2026 and developing tailored protocols that extend Hydra's functionality.

Deliverables

#	Description	Status
D1	Internal report/survey mapping the current existing layer 2 protocols and its relevance/compatibility with Hydra Suite of Protocols	In progress

(iv) Results beyond state of the art

While optimization tools for layer-2 systems have been explored in the context of payment channels and rollups, this work represents the first effort to systematically assess and adapt them to Hydra's unique isomorphic state channel architecture. The internal survey delivers a comparative framework that goes beyond prior general-purpose studies by focusing specifically on Hydra's needs, identifying inefficiencies not addressed in other protocols.

By grounding optimization in Hydra's composability and scalability goals, this research advances the state of the art from generic tools toward ecosystem-specific solutions. The outcome is a blueprint for companion protocols that can sustain throughput, mitigate liquidity bottlenecks, and improve user experience—thereby laying the groundwork for Hydra to evolve into a more robust and versatile layer-2 platform.

IC-3 Light client infrastructure

Start date: Jul 23
Duration (months): 48

(i) Stream overview

This stream addresses a critical gap in the Cardano ecosystem: the absence of a robust, future-proof light client infrastructure. Light clients—designed to provide strong security guarantees without requiring full-node resources—are essential for applications ranging from wallets and DApps to cross-chain protocols like zk-bridges.

Current approaches either rely too heavily on third-party services or lack incentive structures to ensure decentralization and resilience. This project will design light clients that balance security, scalability, and practical constraints such as limited bandwidth, storage, and device capacity. The research will also explore efficient event-notification mechanisms and query frameworks that allow clients to interact with the blockchain reliably while minimizing trust assumptions.

(ii) Context and objectives

As Cardano grows, the size and complexity of its ledger will exceed the capacity of most devices, creating barriers for everyday users and developers. Without secure and decentralized light clients, access to blockchain state risks becoming centralized in third-party services, undermining transparency and trust. This research therefore pursues three core objectives:

1. **Design and formalize secure light client protocols** capable of wallet and smart contract functionality.
2. **Develop standardized query and notification mechanisms** that reduce reliance on full nodes while ensuring fair compensation for service providers.
3. **Establish incentive models** that support adoption and long-term sustainability.

By grounding the work in formal security proofs and composable architectures, the project aims to deliver a reliable foundation for Cardano's future client infrastructure. Success will expand accessibility, reduce reliance on proprietary intermediaries, and ensure that Cardano can support large-scale applications while remaining secure, decentralized, and user-friendly.

(iii) Work performed and main achievements

The core work of this research stream has focused on designing and theoretically analyzing a novel light client protocol tailored for UTxO-based ledgers, including Bitcoin, Cardano, and similar blockchains. This effort resulted in the first deliverable, a paper introducing *Cavefish: Communication-Optimal Light Client Protocol for UTxO Ledgers, Pyrros Chaidos, Aggelos Kiayias, Marc Roeschlin, Polina Vinogradova* (under submission), a two-party protocol between a light client and a service provider.

In Cavefish, the service provider can construct a valid transaction on behalf of the client using only minimal information. When paired with HD wallets, the client needs to share just a single public key and chain code, enabling an ultra-light communication footprint with only two rounds of interaction. The team has already begun exploring potential extensions to the protocol.

Deliverables

#	Description	Documentation
D1	Paper for light client protocol: *Cavefish: Communication-Optimal Light Client Protocol for UTxO Ledgers*	Under submission '25
D2	Paper covering extensions	In progress

(iv) Results beyond state of the art

This work demonstrates the feasibility of fundamentally rethinking light client protocols, moving away from the conventional model where clients mimic full-node synchronization under tight resource constraints. Cavefish offers a radically more efficient approach, significantly reducing storage and bandwidth requirements while maintaining secure and user-friendly blockchain interactions. Given these promising results, the IO Innovation team has launched a new stream to further develop and expand this research direction.

IC-4.1 DApps tokenomics

Start date: Jan 25
Duration (months): 24

(i) Stream overview

Launching a new decentralized application, system, or partner chain involves a critical early phase that often determines long-term success. This stream investigates tokenomic models tailored to this finite "launching phase," where the new entity still depends on a mainchain for security and resources.

Key questions include how to balance partial reliance on mainchain infrastructure, the coexistence of mainchain and partnerchain tokens, the role of initial investment by designers, and the distribution of mutual benefits between both parties. By building new mathematical models that capture these dynamics, the research will provide tools to optimize launch conditions, support sustainable adoption, and ensure a smooth transition toward economic autonomy.

(ii) Context and objectives

Existing tokenomics research largely focuses on steady-state, long-term equilibria, overlooking the challenges of bootstrapping new systems. In early phases, individual users or investors can significantly influence economic outcomes, making assumptions of negligible user impact unrealistic. This research aims to bridge that gap by formalizing launch-phase dynamics and identifying key parameters that drive success—such as adoption thresholds, liquidity requirements, and resource allocations. Objectives include

1. **Defining metrics** for measuring launch success.
2. **Developing models** that quantify trade-offs between designer investment, mainchain support, and participant incentives.
3. **Outlining conditions** under which a DApp or partnerchain can graduate to autonomous, shock-resistant operation.

The ultimate goal is to strengthen Cardano's ability to incubate and integrate new applications and partnerchains, broadening ecosystem growth while maintaining economic stability.

(iii) Work performed and main achievements

In the first half of 2025, research progressed on understanding tokenomics for the launch phase of new systems. Building on the two-token framework developed in TO-1 Tokenomics Design, the team examined how mainchain and partnerchain tokens can coexist, incorporating insights from recent studies on airdrop-based launches and self-reinforcing participation strategies. This included formalizing a game-theoretic model of airdrops, analyzing how equilibria emerge under varying participation costs, token allocations, and network effects. Key results were featured in *Airdrop games*, Sotiris Georganas, Aggelos Kiayias, Paolo Penna, (IJCAI '25), while the paper *Self-Reinforcing Airdrops* (under submission) lays the theoretical foundations for launch-phase token distribution mechanisms and their connection to long-term adoption and incentive alignment.

Beyond theoretical development, this stream has incorporated behavioral and experimental game theory into tokenomics research. A novel survey-based study, *Investors vs Gamblers: Demographics and Attitudes of Cryptocurrency Holders, Sotiris Georganas, Aggelos Kiayias, Charoula Pariarou, Paolo Penna, Alina Velias*, revealed distinctive demographic and attitudinal traits that shape participation and risk preferences. The results were presented at the Conference on Research on Economic Theory and Econometrics (CRETE '25), with a further technical report in progress. These behavioral insights, when combined with the airdrop equilibrium analysis, are shaping a new science of tokenomics for system launches.

Deliverables

#	Description	Status
D1	Paper on key heterogeneity and behavioral aspects capturing fundamental components of ecosystem: *Investors vs gamblers: demographics and attitudes of cryptocurrency holders*	Under submission '25
D2	Paper on general partnerchain and airdrops strategies (a) *Airdrop games* (b) *Self-Reinforcing Airdrops*	(a) Published IJCAI '25 (b) Under submission '25

(iv) Results beyond state of the art

This research shifts tokenomics analysis from long-term equilibria to the underexplored launch phase, where early participation is critical to success. The airdrop game framework, first introduced in academic literature and outlined in the IOG blog post *Airdrop games: towards a theory of tokenomics for blockchain system launches*, shows how incentive design influences whether systems converge on high- or low-engagement equilibria. Beyond identifying equilibria, it assesses their likelihood, giving designers practical levers—such as targeted allocations or lowering entry costs—to promote sustainable growth. This moves well beyond prior static models that assume negligible individual influence on outcomes.

The integration of behavioral insights is another key advance. By combining mathematical rigor with empirical findings on crypto user attitudes, this work bridges economics, cryptography, and behavioral science. It marks the first systematic effort to align launch mechanisms with the observed characteristics of crypto holders, rather than assuming purely rational actors. This interdisciplinary innovation positions Cardano at the frontier of tokenomics research, with immediate application to partnerchains and DApps, while contributing broadly to the design of resilient and sustainable blockchain economies.

IC-4.2 Consensus innovation

Start date: Sep 24
Duration (months): 54

(i) Project Summary

This research stream explores next-generation consensus protocols for Cardano, advancing beyond current Ouroboros models by integrating insights from cryptography and distributed systems. It focuses on extending both Nakamoto-style probabilistic consensus and BFT-style protocols, while also investigating emerging paradigms such as DAG-based consensus. By relaxing synchronous assumptions, adapting to dynamic participation, and improving efficiency, the work aims to design consensus mechanisms that deliver stronger security, scalability, and decentralization. The outputs will support Cardano's PartnerChains framework and, in the long term, inform improvements to Cardano's backbone consensus itself.

(ii) Context and objectives

While state-of-the-art consensus protocols like HotStuff, Jolteon, and Narwhal and Tusk have made significant progress, most assume static validator sets and operate in semi-synchronous environments, limiting their applicability to permissionless proof-of-stake systems. This research addresses those gaps:

1. **Developing competitive algorithms** tailored for Cardano's PartnerChains.
2. **Exploring ways to improve Ouroboros** from a security and operational perspective
3. **Exploring future designs such as DAG-based algorithms -** to foster excellence in the design of dynamically available PoS protocols.

The overarching objective is to establish a portfolio of consensus solutions that push the boundaries of performance and robustness across various core designs, with measurable improvements in security, latency, and throughput, and adoption in at least one or two flagship projects.

(iii) Work performed and main achievements

In the first half of 2025, progress was made along two major research directions. On the protocol side, we advanced the formal verification of iterated BFT-style consensus, focusing specifically on *Jolteon*. Mechanized proofs for safety have been completed, with liveness proofs well underway. These results are available in a public GitHub repository, ensuring transparency and accessibility to the community. Work is ongoing to extend this analysis to settings with stake-shift, thereby capturing more realistic dynamics of proof-of-stake systems.

In parallel, a new game-theoretic model was developed to explain why actors in PoS systems are incentivized to mint blocks on time, with emphasis on whether Nakamoto-style consensus needs additional mechanisms for greater economic robustness. A first draft is complete, with a research paper expected by year-end. In addition, anti-grinding has been thoroughly investigated—both generally and for Ouroboros—resulting in a dedicated paper *Taming Iterative Grinding Attacks on Blockchain Beacons*, *Peter Gaži, Saad Quader, Alex Russell* accepted at Asiacrypt '25. Together, these advances strengthen the theoretical and practical foundations of Cardano's consensus research.

Deliverables

#	Description	Status
D1	Paper for provably secure iterated BFT protocols including mechanized proofs	In progress
D2	Paper on incentives in Ouroboros	In progress
D3	Paper on anti-grinding: *Taming Iterative Grinding Attacks on Blockchain Beacons"*	Accepted Asiacrypt '25

(iv) Results beyond state of the art

This research stream pushes beyond the state of the art in two critical dimensions. First, by delivering mechanized safety and liveness proofs for *Jolteon*, it elevates consensus analysis from informal reasoning and pen-and-paper proofs to fully mechanized, machine-checked verification. This significantly raises the assurance level compared to existing consensus literature, where such mechanization is rare, especially for complex proof-of-stake protocols. The planned extension to stake-shift analysis represents another novel step, addressing a real-world factor largely absent in current formal treatments.

Second, the incentive modeling introduces a rigorous game-theoretic framework for PoS systems that explains why block producers act honestly under varying protocol rules and market conditions, bridging consensus guarantees with incentive compatibility for more robust, economically sustainable mechanisms. In parallel, anti-grinding measures were advanced to limit adversarial influence over randomness. Together, these results position Cardano at the forefront of consensus innovation, shaping both research and next-generation protocol design.

7. Change requests

Given the exploratory nature of early-stage research, flexibility is essential. IOR manages its work through a portfolio-based approach, allowing resources to be directed toward the most promising and strategically aligned streams. This enables regular reassessment of progress, risks, and opportunities, with strong projects accelerated and less aligned ones scaled back.

With this in mind, IOR proposes to deprioritise the following three streams:

- *WOS-2: State-Machine Contract Environment* – delayed by around six months due to hiring challenges. While the work remains highly valuable—developing a higher-level, formally verifiable framework for EUTxO smart contracts, including Hydra contracts—it is recommended for WP26 rescheduling.

- *IHT-4: Hydra Auditing Tools* – reducing Hydra research from four to three streams in response to community feedback. Auditing will be revisited once further L2 capabilities funded in WP25 are online and use cases are clearer.

- *GI-1: Global Identity* – resources have been redirected toward Blockchain DeFi, leaving only the formalization of Universal Anonymous Signatures v2 in scope. This next version, enabling more flexible privacy-preserving identities, is expected in 2026 and the stream will be reviewed then.

IOR therefore proposes to prioritize three streams: *WOS-4 Decentralized Storage*, as results are being consolidated into a report and potential paper; *IC-1 State Proofs & Bridges*, with active work across core research, ZK, and TEE strands; and *ZK-1 Core zero-knowledge capabilities*, a cross-cutting stream delivering notable outputs across multiple streams.

WOS-4: Decentralized Storage

Start date: Jan 25
Duration (months): 24

(i) Stream overview

This research focuses on creating a Byzantine-resilient storage protocol that withstands malicious or corrupted nodes in permissionless environments. The goal is a provably secure system that ensures integrity, persistence, and retrievability, either by strengthening existing Distributed Hash Tables (DHTs) or introducing a new approach.

Such a protocol would extend Cardano's capabilities beyond UTxOs—enabling NFT file persistence, decentralized pub/sub systems, and smart contracts with direct data access. Optional extensions include smart contract integration and an incentive layer to reward nodes for storage and retrieval, ensuring broad participation.

(ii) Context and objectives

Current decentralized storage solutions like IPFS and Filecoin demonstrate potential but face limitations in adversarial, permissionless settings. Building on concepts such as proofs of storage, replication, and retrievability, this project seeks a Cardano-specific solution with security guarantees equivalent to mainchain consensus. The objectives are to:

1. **Design a Byzantine-resilient storage protocol** with formal guarantees of security, availability, and retrievability, even under adaptive attacks.
2. **Prototype and evaluate performance** across efficiency, redundancy, resource use, and adversarial robustness.
3. **Extend smart contract functionality** to enable secure storage and retrieval of external data (e.g., NFT-linked files, off-chain state).
4. **Develop an incentive layer** to reward reliable storage nodes, with higher rewards for underrepresented or high-risk regions *(Optional)*.

Delivering on these goals will provide Cardano with a secure persistence layer that unlocks new decentralized applications, improves resilience, and strengthens its position as a leading Web3 platform.

(iii) Work performed and main achievements

In 2025, this workstream focused on two fronts: a formal treatment of adaptively secure decentralized storage networks (DSNs), and the development of Byzantine-resilient primitives for Distributed Hash Tables (DHTs) and Data Availability Sampling (DAS).

This first formalizes vulnerabilities in proof-of-replication (PoRep) when combined with blockchains, introducing a framework for adaptive attacks during storage and retrieval. It proposes a generic DSN construction over EUTxO with automatic, client-free PoRep verification, and proves that succinctness and non-malleability are essential for security—showing experimentally that without them, systems are vulnerable to attacks like offloading data while still producing valid proofs.

For DHTs, research focused on Sybil- and Byzantine-resistant designs, evaluating proof-of-space as a defense and exploring whether proof-of-stake could offer similar guarantees under shifting stake distributions. In parallel, DAS research examined whether its separation from Robust Distributed Arrays is necessary, or if a unified primitive could improve efficiency without weakening availability. Achievements include clarifying research directions around PoSpace vs. PoStake defenses and identifying pathways to simplify DAS and RDA, with results to be consolidated in a report and potentially a research paper later in 2025.

Deliverables

#	Description	Status
D1	Paper defining the security model: *Adaptively Secure Blockchain-Aided Decentralized Storage Networks: Formalization and Generic Construction*	Under submission '25
D2	Paper on presenting the protocol design and providing a formal security proof	In progress

(iv) Results beyond state of the art

This workstream extends beyond existing decentralized storage and availability solutions, such as IPFS and Filecoin, which rely heavily on heuristic defenses and lack robust guarantees in adversarial settings. By investigating formally grounded mechanisms for Sybil resistance in DHTs—specifically the use of PoSpace and potentially PoStake—the project seeks to provide provable security even under adaptive adversaries. This represents a significant step forward from today's systems, where adaptive strategies can often undermine protocol assumptions.

In the DAS domain, the exploration of whether Data Availability Sampling and Robust Distributed Arrays can be unified into a single primitive also pushes beyond the current state of the art. Such a unification could lead to more efficient protocols, lower complexity in implementation, and stronger availability guarantees. Together, these directions aim to replace heuristic approaches with provably secure, Byzantine-resilient primitives, laying the foundation for a decentralized storage layer on Cardano that is both efficient and formally verifiable.

IC-1: State proofs and blockchain bridges

Start date: Mar 24
Duration (months): 36

(i) Stream overview

This research stream investigates the design of secure, efficient, and trust-minimized blockchain bridges, with a primary focus on enabling cross-ledger interoperability between Cardano and Midnight. Unlike committee-based bridges, which rely on trusted operators and have suffered from high-profile security breaches, this work emphasizes trustless designs that eliminate single points of failure.

This stream studies the use of state proofs, consensus certificates, and zero-knowledge tooling to transfer assets, tokens, and arbitrary state across chains in a secure and scalable manner. In parallel, it is evaluating complementary approaches such as committee-based designs and Trusted Execution Environments (TEEs), balancing performance and security. The resulting bridge protocols will serve as foundational infrastructure for cross-chain applications, enabling asset transfer, contract portability, and broader ecosystem integration.

(ii) Context and objectives

Blockchain bridges are critical for ecosystem growth, yet most existing implementations face trade-offs between efficiency, security, and decentralization. This research aims to

1. **Formally analyze the security and functionality** of bridge protocols in both game-based and composable (UC) frameworks.
2. **Design and optimize zero-knowledge circuits** for proving state transitions and consensus validity, ensuring practical performance for real-world deployment.
3. **Investigate TEEs** as an alternative or hybrid approach, simplifying bridge design while measuring their efficiency relative to cryptographic solutions.

The key objectives are to deliver a provably secure protocol for asset and information transfer for Partner Chains, extend these designs to more general cross-chain state transfers, and establish benchmarks for efficiency, scalability, and privacy. Success will be measured through reductions in overhead relative to centralized bridges, adoption, and enablement of novel cross-ledger.

(iii) Work performed and main achievements

The core research strand achieved a major milestone with the paper *Bridge Games: A Formal Treatment of Security for Ledgers, Bridges, and Cross-Chain Applications, Pyrros Chaidos, Pooya Farshim, Dimitar Jetchev, Aggelos Kiayias, Markulf Kohlweiss, Anca Nitulescu* (under submission). This work lays the first rigorous security definitions for bridges, explores different constructions under varying cryptographic and trust assumptions, and analyzes a basic cross-chain application (token transfer). Future work may extend this by analyzing real-world bridges such as Polygon or Polkadot as well as the recent BitVM bitcoin bridge.

The ZK strand advanced efficiency for trustless bridges, with the paper *Efficient Batch Opening Schemes for Merkle Tree Commitments with Applications to Trustless Crosschain Bridge*, Bingsheng Zhang, Wuyunsiqin, Xun Zhang, Markulf Kohlweiss, Kui Ren (ICCCN '25). These schemes enable succinct proofs of multiple data points, lowering on-chain verification costs. Supporting software artifacts are in development to provide realistic performance benchmarks.

The TEE strand investigates transparent TEEs, which may leak randomness or inputs, and studies how to build UC-secure zkSNARKs using them as SNARK replacements. This work supports the ZK-1 D3 TEE-Based ZK proof servers paper (see below), it is closely aligned with a TEE Bridge innovation stream, and a dedicated paper is planned for late 2025.

Deliverables

#	Deliverable	Status
D1	Paper on foundational security definitions for bridges: *Bridge Games: A Formal Treatment of Security for Ledgers, Bridges, and Cross-Chain Applications*	Under submission '25
D2	Paper on ZK proofs: *Efficient Batch Opening Schemes for Merkle Tree Commitment with Applications to Trustless Crosschain Bridge*	Published ICCCN '25

(iv) Results beyond the state of the art

This work delivers the first formal security framework for blockchain bridges, defining clear, code-based notions of soundness and liveness within the Universal Composition (UC) model. It enables modular bridge designs that can adapt to various trust and blockchain environments by expressing protocols as interchangeable components.

The research categorizes popular bridge types—such as direct-observation, certificate-based, and rotating-committee designs—and proposes new cryptographic constructions using threshold multi-signatures. It also unifies light clients and zkBridges under the same model, enabling consistent and scalable cross-chain design.

Going beyond the state of the art, this research replaces ad-hoc designs with a rigorous foundation for building, analyzing, and improving bridges. For Cardano, it paves the way for secure interoperability, supporting safe cross-chain transactions and expanding the network's reach.

ZK-1 Core zero-knowledge capabilities

Start date: Jan 24
Duration (months): 48

(i) Stream overview

This stream establishes a unified technical foundation for zero-knowledge (ZK) capabilities across the Cardano ecosystem, supporting use cases such as light clients, state proofs, blockchain bridges, and Hydra Tail. Given the rapid pace of advancement in ZK technology and its growing importance both internally and across the blockchain industry, Cardano must adopt a modular, upgradable approach to ZK tooling.

This includes active participation in standardization efforts—such as the Halo2 proving system—to ensure interoperability with the broader ZK research and engineering community. By consolidating expertise and aligning around shared primitives, this work will enable sustainable, future-proof ZK infrastructure for Cardano's ecosystem.

(ii) Context and objectives

Zero-knowledge proofs are becoming critical for scalability, privacy, and verifiability in decentralized systems, yet they remain a specialized and evolving domain where research and engineering intersect. This project aims to:

1. **Ensure Cardano's ZK stack is modular** and easily upgradeable, allowing seamless integration of new proof systems.
2. **Expand ZK support in Plutus** smart contracts, starting with state proofs and progressing toward arbitrary contract logic.
3. **Explore advanced techniques** like recursive proofs and folding schemes to enable efficient verification of complex computations on-chain.
4. **Foster stronger ties** with the global ZK community through joint research and standardization.

The outcome will be a robust, evidence-based ZK toolkit underpinning a wide range of applications—from governance and financial services to data provenance and gaming—while ensuring Cardano remains at the forefront of ZK innovation.

(iii) Work performed and main achievements

In the first half of 2025, this stream has focused on building foundational capacity in zero-knowledge (ZK) techniques for Cardano, coordinating across related workstreams and research groups. A major milestone was the advancement of universally composable (UC) proof systems. At CSF '25, the team presented *AGATE: Augmented Global Attested Trusted Execution in the Universal Composability Framework*, *Lorenzo Martinico, Markulf Kohlweiss,* which introduced a modular UC definition of Trusted Execution Environments (TEEs). This work enables meaningful comparisons between different enclave capabilities and adversarial models, and shows how weaker TEEs can be composed to emulate stronger ones. Ongoing research uses TEEs as opposed to global generic groups, with plans to model proof servers in a paper on TEE-based ZK proof servers.

In parallel, earlier progress on UC security for succinct proofs was consolidated through *The Brave New World of Global Generic Groups and UC-Secure Zero-Overhead SNARKs*, Jan *Bobolz, Pooya Farshim, Markulf Kohlweiss, Akira Takahashi* (TCC '24), which demonstrated that Groth16—a widely used SNARK—can be proven UC-secure without efficiency loss. This work was also presented at the ZKProof7 workshop in Sofia, and together these results strengthen the theoretical foundation for Cardano's future ZK infrastructure. *Universally Composable SNARKs with Transparent Setup without Programmable Random Oracle*, *Christian Badertscher , Matteo Campanelli, Michele Ciampi, Luigi Russo, Luisa Siniscalchi*, overcomes long-standing assumptions in the field and introduces new modeling techniques and cryptographic tools that make secure and composable zero-knowledge proofs more practical and broadly applicable.

Beyond theory, the stream contributed to applied research. Work on enhancing Mithril signatures with ZK techniques has delivered progress on more succinct aggregation, particularly through Boneh–Lynn–Shacham signature aggregation and Merkle tree inclusion proofs. At the same time, new directions in post-quantum cryptography have been initiated and investigations into lattice-based SNARKs and folding schemes. These efforts produced an early Systematization of Knowledge (SoK) survey on lattice-based folding, polynomial commitments, and SNARK constructions, setting the stage for a roadmap in post-quantum ZK research.

Deliverables

#	Deliverable	Status
D2	Paper on TEE environments: *AGATE: Augmented Global Attested Trusted Execution in the Universal Composability Framework*	Published CSF '25
D3	Paper on TEE-based ZK proof servers: *"Brave New World of Observable Programmable Trusted Execution Environments and Zero-Overhead SNARKs"*	In progress
D4	Paper on SNARKs with minimal setup: *Universally Composable SNARKs with Transparent Setup without Programmable Random Oracle*	Published Crypto '25
D5	Paper on Systemization of Knowledge: Lattice-based Folding and Polynomial Commitment Schemes	In progress

(iv) Results beyond the state of the art

This stream pushes beyond the state of the art by combining rigorous UC frameworks with practical blockchain applications. The AGATE framework represents a significant advance over prior models like G-ATT, introducing a modular treatment of TEEs that captures real-world diversity in enclave capabilities. It equips protocol designers with tools to reason precisely about what security guarantees different TEEs can (and cannot) provide, enabling stronger integration of hardware-assisted trust into ZK protocols. Complementing this, the UC-secure treatment of Groth16 breaks new ground by proving composable security without overhead—bridging a long-standing divide between standalone and UC-secure SNARKs and setting a precedent for efficient, rigorously composable proof systems.

Applied results also extend the frontier. The integration of ZK into Mithril signature aggregation introduces a new level of efficiency for lightweight clients and cross-chain verification, far beyond what is available in existing certificate aggregation mechanisms. Meanwhile, the exploration of lattice-based SNARKs and folding schemes addresses the urgent challenge of building quantum-resistant proofs, a direction still underexplored in the ZK space. By coupling immediate ecosystem benefits—such as efficient proof aggregation—with long-term resilience through post-quantum research, this stream ensures that Cardano remains at the forefront of zero-knowledge innovation.

8. Communication and dissemination

IOR actively shares its research through company channels, community focused events, and academic publications, fostering transparency, collaboration, and impact across the Cardano community and beyond.

IOHK channels

The IOHK website hosts its comprehensive library of papers and curated YouTube content, further enhanced through video campaigns and playlists showcasing researcher insights, protocol advancements, and academic partnerships on the IOHK YouTube channel.

The IOHK blog highlights IO Research's groundbreaking contributions to blockchain innovation, featuring in-depth explorations of key advancements, collaborations, and thought leadership. Notable posts include:

- Professor Kiayias explores consensus evolution at Science of Blockchain Conference 2025 - insights on past developments and upcoming challenges for securing decentralized systems
- A new era of smart contract verification on Cardano - introducing a powerful new tool designed to make smart contracts faster, simpler, and more accessible
- Input | Output co-hosts 'Cryptographic tools for blockchains' workshop at Eurocrypt '25 - bringing together leading researchers to explore and advance cutting-edge cryptographic techniques for distributed ledger technologies
- Leios takes the stage at Crypto 2025 - the IOR presented its groundbreaking analysis of Ouroboros Leios to significantly boost Cardano's scalability and throughput.
- Airdrop games: towards a theory of tokenomics for blockchain system launches - a game-theoretic model providing designers with tools to strategically allocate tokens and guide user participation (as per above)
- IOR outlines a research framework for SPOs - outlining focused research priorities aimed at enhancing infrastructure security, scalability, and operational resilience
- Ouroboros Peras: accelerating transaction settlement on Cardano - speeding up transaction settlement to around two minutes using a voting-based chain selection

IOR also runs micro campaigns on social media that support its ecosystem, partnerships, dissemination and other strategic objectives. Notable campaigns included:

- Aggelos Kiyias explains at SBC '25 *(3k views)* that decentralization is a foundation for global systems that are fair, transparent, and resilient at SBC '25
- Paolo Penna highlighted three tokenomics streams *(32k views)* and stresses the need for integrated models that view Cardano as a living economy
- Outlining the Hydra research direction *(40k views)* for the broader L2 ecosystem is essential to supporting the community scale Cardano
- Fergie Miller interviews with BigPey *(6k views)* to outline the IOR vision for the future of Cardano to 2030

Cardano R&D Sessions

IOR revamped the Research Working Group (RWG) monthly meetings into monthly R&D Sessions for the Cardano community. Each session features thematic discussions with invited guest speakers from relevant community segments, planned quarterly. So far, three sessions have been run with further sessions planned for this year including Ourborous, Interchains and Worlds'operating System

- Layer 2 Expansion - Beyond Hydra, June - The inaugural session highlighted the growing diversity of L2 solutions in the ecosystem. Sandro Coretti-Drayton (IOR) outlined Hydra's latest advances—including Heads, Tails, inter-head networking, and auditing tools—while stressing its role as one piece of a broader scaling vision.

  A panel with Midgard, zkFold, Eryx, and Gummiworm explored optimistic rollups, zk-rollups, ZK bridging, and scalable rollup networks, focusing on interoperability, efficiency, and shared standards. Emphasizing security, sustainability, and collaboration, the session underscored how Cardano's EUTXO model supports a modular, pluralistic L2 ecosystem.

- Cardano Tokenomics, July - The July session explored tokenomics as a driver of sustainable growth, decentralization, and real-world utility. Paolo Penna presented models of validator incentives, pledge mechanics, and stablecoin design, while Ryan Wiley introduced CIP-50 to strengthen decentralization via revised pledge penalties.

  A panel with the Cardano Foundation and Fluid Tokens emphasized transparency, interoperability, and stablecoins as key to adoption. With Cardano moving toward a self-sustaining economic model, the session highlighted that effective tokenomics is not just parameter tuning but the foundation of trust, participation, and resilience.

- Technology Validation, August - The IO R&D Innovation team translates advanced research into practical solutions using an evidence-based approach of formal methods, simulations, and prototyping.

  Key highlights include: Phalanx, making certain consensus attacks billions of times harder to improve finality; Jolteon, with formally proven liveness for stronger partner chain security; RSnarks, enabling efficient on-chain Halo2 proof verification for trustless bridges and rollups; Minotaur, a design for new blockchains to re-stake security from Cardano or Ethereum; Cavefish, a lightweight client for UTXO chains without full ledger sync; and Committee Proofs (proposed), to enable secure cross-chain bridges with rotating committees.

Research conferences

In the first half of 2025, IOR presented the research outlined in this report at several major international conferences, contributing to the advancement of blockchain science. These venues provided a platform to share findings, engage with the academic community, and advance the goals of Cardano Vision. The table below lists the conferences attended, with links to each event:

Short conference name	Full name	Link to event
FC ‘25	International Conference on Financial Cryptography and Data Security -	https://fc25.ifca.ai/
FMBC '25	International Workshop on Formal Methods for Blockchains	https://fmbc.gitlab.io/2025/
CSF '25	IEEE Computer Security Foundations Symposium	https://csf2025.ieee-security.org/
IJCAI '25	International Joint Conference on Artificial Intelligence	https://www.ijcai.org/
ITC '25	Conference on Information-Theoretic Cryptography	https://itcrypto.github.io/2025/
Crypto '25	International Cryptology Conference	https://crypto.iacr.org/

Appendix

A. Proposal documentation

The Input Output Research (IOR): Cardano Vision - Work Program 2025 proposal consists of the following documents:

- Input Output Research (IOR): Cardano Vision - Work Program 2025 Proposal v1.0.pdf
- Input Output Research (IOR): Work Program 2025 - Proposed Fundamental Research Streams v1.0
- Input Output Research (IOR): Work Program 2025 - Proposed Technology Validation Streams v1.0
- 25.04 Intersect - Product Committee - Research Working Group - Cardano Vision v1.06

B. Table of deliverables

#	Stream ID	Paper	Status
1	WOS-4	*Adaptively Secure Blockchain-Aided Decentralized Storage Networks: Formalization and Generic Construction*, Xiangyu Su, Yuma Tamagawa, Mario Larangeira, Keisuke Tanaka	Under submission '25
2	WOS-6	*Incentivizing Geographic Diversity for Decentralized Systems*, Raghav Bhaskar, Aggelos Kiayias, Evangelos Markakis, Marc Roeschlin	Under submission '25
3	OO-1V	*Adaptively Secure Fast Settlement with Dynamic Participation and Self-Healing*, Christian Badertscher, Sandro Coretti-Drayton, Peter Gaži, Aggelos Kiayias, Alexander Russell	In progress - expected 2025
4	OO-2	*High-Throughput Permissionless Blockchain Consensus under*	Published Crypto '25

		Realistic Network Assumptions, Sandro Coretti-Drayton, Matthias Fitzi, Aggelos Kiayias, Giorgos Panagiotakos, Alexander Russell	
5	OO-3	*Universally Composable Transaction Order Fairness*: Refined Definitions and Adaptive Security, Michele Ciampi, Aggelos Kiayias, Yu Shen (published shortly)	Accepted Asiacrypt '25
6	OO-5	Paper describing a provably secure implementation of multi-resource BFT consensus or its components	In progress - expected 2025
7	OO-6	*Proof-of-Useful-Work is Practical: Consensus via a Competitive Optimization Engine*, Matthias Fitzi, Aggelos Kiayias, Laurent Michel, Giorgos Panagiotakos, Alexander Russell	Under submission '25
8	OO-7	*One-dimensional vs. Multi-dimensional Pricing in Blockchain Protocols*, Aggelos Kiayias, Elias Koutsoupias, Giorgos Panagiotakos, Kyriaki Zioga	Under submission '25
9	TO-1	*Single-token vs Two-token Blockchain Tokenomics*, Aggelos Kiayias, Philip Lazos, Paolo Penna	Accepted AFT '25
10	TO-1	*Algorithmic Monetary Policies for Blockchain Participation Games*, Diodato Ferraioli, Manvir Schneider, Paolo Penna	Under submission '25
11	TO-2	*Pool Formation in Oceanic Games: Shapley Value and Proportional Sharing*, Aggelos Kiayias, Elias Koutsoupias, Evangelos Markakis, Panagiotis Tsamopoulos	Accepted AFT '25
12	D4-1	*Beyond Blockchain Ballots: UC-Secure Layer-2 Voting and Governance*, Raghav Bhaskar, Pooya Farshim, Matthias Fitzi, Aggelos Kiayias	Under submission '25

13	D4-2	*Reward Schemes and Committee Sizes in Proof of Stake Governance*, Georgios Birmpas, Philip Lazos, Evangelos Markakis, Paolo Penna	Published FC '25
14	D4-2	Paper on participatory budgeting	In progress - expected 2025
15	IHT-1	Paper on security proofs	In progress - expected 2025
16	IHT-2	*Universally Composable Treatment of Multi-Party Isomorphic State Channels*, Maxim Jourenko, Xiangyu Su, Adam Blatchley Hansen, Mario Larangeira	Under submission '25
17	IC-1	*Bridge Games: A Formal Treatment of Security for Ledgers, Bridges, and Cross-Chain Applications. Pyrros Chaidos, Pooya Farshim, Dimitar Jetchev, Aggelos Kiayias, Markulf Kohlweiss, Anca Nitulescu*	Under submission '25
18	IC-1	*Efficient Batch Opening Schemes for Merkle Tree Commitment with Applications to Trustless Crosschain Bridge*, Bingsheng Zhang, Wuyunsiqin, Xun Zhang, Markulf Kohlweiss, Kui Ren	Published ICCCN '25
19	IC-3	*Cavefish: Communication-Optimal Light Client Protocol for UTxO Ledgers*, Pyrros Chaidos, Aggelos Kiayias, Marc Roeschlin, Polina Vinogradova	Under submission '25
20	IC4.1	*Investors vs gamblers: demographics and attitudes of cryptocurrency holders*, Sotiris Georganas, Aggelos Kiayias, Charoula Pariarou, Paolo Penna, Alina Velias.	Under submission '25
21	IC4.1	*Airdrop games*, Sotiris Georganas, Aggelos Kiayias, Paolo Penna	Published IJCAI '25

22	IC4-1	*Self-Reinforcing Airdrops: Less Giveaway More Participation*, Sotiris Georganas, Aggelos Kiayias, Paolo Penna	Under submission ‘25
23	IC4.2	*Taming Iterative Grinding Attacks on Blockchain Beacons*, Peter Gaži, Saad Quader, Alex Russell	Accepted Asiacrypt ‘25
24	IC4.2	Paper on incentives in Ouroboros	In progress - expected 2025
25	ZK-1	*AGATE: Augmented Global Attested Trusted Execution in the Universal Composability Framework, Lorenzo Martinico, Markulf Kohlweiss*	Published CSF ‘25
26	ZK-1	Paper on TEE-Based ZK proof server: *Brave New World of Observable Programmable Trusted Execution Environments and Zero-Overhead SNARKs*	In progress
27	ZK-1	*Universally Composable SNARKs with Transparent Setup without Programmable Random Oracle, Christian Badertscher, Matteo Campanelli, Michele Ciampi, Luigi Russo, and Luisa Siniscalchi*	Published Crypto ‘25
28	ZK-1	Paper on Systemization of Knowledge: *Lattice-based Folding and Polynomial Commitment Schemes*	In progress

In addition to the research outputs above, the following relevant paper and reports will be provided in 2025:

WOS-3: *A Layered Certifying Compiler Architecture*, Jacco Krijnen, Joris Dral, Manuel Chakravarty, Gabriele Keller, Wouter Swierstra, accepted at FUNARCH '25

This paper proposes a practical approach to ensuring compiler correctness using a *layered certifying architecture*. Instead of verifying the compiler and its correctness proof together—a method ill-suited for rapidly evolving open-source projects—it introduces a separate certifier that applies *translation validation* to verify each compiler run. By structuring the certifier in independent, incremental layers, it remains decoupled from the main compiler's frequent changes. The approach is demonstrated in production through a certifier for the Plutus smart contract compiler, with functional languages and proof assistants like Rocq enabling seamless integration and proof extraction.

OO-6: Applicability report of PoUW to ML domain

This report explores *Proof-of-Useful-Work (PoUW)* for machine learning, focusing on *Proof-of-Deep-Learning (PoDL)*—a system that replaces energy-wasting cryptographic mining with useful deep learning model training. To address inefficiencies in training, the framework includes a Transformer-based predictor and a stabilization detection algorithm that identifies when model loss saturates, allowing for early termination of unproductive training.

IC4-1: A follow up technical report to *Investors vs gamblers: demographics and attitudes of cryptocurrency holders*, Georganas et al., 2025

A forthcoming technical report will present expanded findings from a large-scale, multi-country survey of up to 11,000 participants across the USA, UK, Germany, and Greece. Results indicate that 15–20% of respondents hold crypto assets, with crypto holders generally more risk-averse and patient than stock investors, though with country-specific variations. Incorporating additional survey data, the report will refine and extend a mathematical model to better capture these behavioural patterns within existing economic frameworks.

C. Change request project plans

WOS-4: Decentralized storage

Research Lead: Sandro Coretti-Drayton
Start date: Jan 24
Forecast Duration (months): 24
2025 FTEs: 1.4
Intersect Product Roadmap: Architectural Excellence

Objectives	The primary objective of this research is to develop a distributed file storage protocol that is both robust and practical. It must be resilient to Byzantine faults and provably secure under assumptions similar to those of the Cardano mainchain. A key challenge is defending against adaptive attackers, who can target nodes based on protocol activity. Since only a subset of nodes will store any given file, this creates a vulnerability to targeted corruption or denial-of-service (DoS) attacks.
Workplan	T1. Define the security model, encompassing the network structure, characteristics of the adversary, and underlying assumptions. (Start: M1, Duration: 4 months) T2. Settle on security definition in the model (property-based or UC). (Start: M1, Duration: 4 months) T3. Design a distributed storage protocol that aligns with the established security model. (Start: M5, Duration: 18 months) T4. Prove the security of the proposed protocol w.r.t. to the above security definition. (Start: M5, Duration: 18 months) T5. Write a research paper. (Start: M21, Duration: 4 months)
Deliverables	D1. A paper defining the security model, presenting the protocol design, and providing a formal security proof.
Impact	Enables robust persistence for NFT files and off-chain data, unlocking new DApp capabilities, decentralized pub/sub systems, and smart contracts with direct data access. Enhances network resilience through a formally verifiable storage layer, expanding core ecosystem utility and functionality for developers.
Minimal Requirements	n/a

Team	Research Fellow, 0.25 FTE (24 months) Research Fellow, 0.25 FTE (24 months) Researcher, 0.25 FTE (24 months) Chief Scientist, 0.1 FTE (24 months)

IC-1: State proofs and blockchain bridges

Research Lead: Pooya Farshim
Start date: Mar 24
Forecast Duration (months): 36
2025 FTEs: 1.9
Intersect Product Roadmap: Architectural Excellence

Objectives	The goal of the project is to study trustless bridges from a first principle perspective, paying particular attention to different design approaches and their underlying trust assumptions.
Workplan	T1. Core Strand: Define and analyze secure cross-chain communication, comparing trustless and trusted bridges, and specifying protocol components, primitives, and safeguards. (Start: M1, Duration: 24 months) T2. ZK Strand: Design ZK circuits for Cardano and Midnight use cases and identifies optimal proving methods for light client compatibility. (Start: M1, Duration: 24 months) T3. TEE Stand: Explore implementing bridge primitives using TEEs with minimal trust, including transparent or memory-limited TEEs, and enhancing them with consensus-based persistent storage. (Start: M1, Duration: 24 months)
Deliverables	D1. Bridge Games: A Formal Treatment of Security for Ledgers, Bridges, and Cross-Chain Applications D2. Efficient Batch Opening Schemes for Merkle Tree Commitment with Applications to Trustless Crosschain Bridge D3. A paper for the TEE strand
Impact	TBC

Minimal Requirements	The bridge protocol must enable, at a minimum, the secure transfer of assets and information from a partner chain to Cardano, both of which are UTxO based.
Team	1. Core research strand: Research Fellow, 0.2 FTE (24 Months) Research Fellow, 0.2 FTE (24 Months) Research Fellow, 0.2 FTE (24 Months) Applied Cryptography Researcher, 0.1 FTE (24 Months) Chief Scientist, 0.1 FTE (24 Months) 2. ZK strand: Research Fellow, 0.1 FTE (24 Months) Research Fellow, 0.2 FTE (24 Months) Research Fellow, 0.3 FTE (24 Months) External, 0.4 FTE (24 Months) 3. TEE strand: External, 0.2 FTE (24 Months) Research Fellow, 0.2 FTE (24 Months) Research Fellow, 0.2 FTE (24 Months) External, 0.2 FTE (24 Months) Research Fellow, 0.2 FTE (24 Months)

ZK-1 Core zero-knowledge capabilities

Research Lead: Markulf Kohlweiss
Start date: Jan 24
Forecast Duration (months): 48
2025 FTEs: 1
Intersect Product Roadmap: Architectural Excellence

Objectives	Ensuring the long-term sustainability and upgradability of Cardano's cryptographic algorithms is essential, especially as zero-knowledge (ZK) protocols continue to evolve rapidly. To remain current and secure, Cardano's ZK tooling must be modular and easily updateable, enabling seamless integration of future advancements such as post-quantum security.

	Other goals include expanding ZK capabilities within Plutus smart contracts—initially for specific use cases like state proofs, and eventually for arbitrary contract logic. A major future milestone is enabling proofs of arbitrarily long computations through recursive proofs and folding schemes, allowing efficient and verifiable computation on-chain.
Workplan	T1. Advancement of universally composable (UC) proof systems T2. Analyse different enclave capabilities and adversarial models, T3. Investigate building scalable, verifiable systems by integrating TEEs T4. Enhancing Mithril signatures with succinct ZK proof techniques T5. Systematization of Knowledge (SoK) survey on cryptographic techniques that t enable efficient, secure zero-knowledge proofs
Deliverables	D1. Paper on universally composable (UC) proof systems D2. Paper on TEE environments D3. Paper on TEE-based ZK proof servers D4. Paper on SNARKs with minimal setup D5. Paper on Systemization of Knowledge
Impact	To build shared, evidence-based zero-knowledge capabilities across the Cardano ecosystem. Integrating ZK proofs into Plutus would significantly enhance smart contract functionality.
Minimal Requirements	n/a
Team	Research Fellow, 0.2 FTE (48 Months) Research Fellow, 0.2 FTE (24 Months) Researcher, 0.2 FTE (24 Months) Researcher, 0.2 FTE (48 Months) External, 0.2 FTE (48 Months)