Resonant Network - Research Overview

An R64N Research Project

Abstract

Human decision-making underpins nearly every domain of society, yet existing empirical methods such as periodic surveys, controlled experiments, and observational datasets each capture only a partial view of how populations make decisions, and none operates as a continuously evolving observatory of structured decisions across domains. Resonant Network implements such an observatory, and the framework described here is operational, with a growing base of registered users and recorded decision events; a longitudinal system in which participants voluntarily complete structured trade-off decisions spanning everyday choices, ethical dilemmas, strategic interaction, cognitive dilemmas, and other domains. Each completed scenario yields a structured decision event. Basic decisions map to an auditable six-dimensional coordinate system whose axes are based on a classification of decision types whereas complex decisions are analyzed topologically. All outputs are de-identified aggregate metadata computed from individual-level data, comprising auditable population statistics alongside non-auditable derived representations such as embeddings and evolving topologies.

Unlike single-session user interactions, decision scenarios are distributed through time, allowing decisions to be observed under changing personal and societal contexts. Participants simultaneously construct a longitudinal personal decision profile that remains under their control, aligning individual utility with scientific observation. While individual scenarios draw upon established survey and vignette methodologies, the contribution of this work lies in their continuous operation, common representation, and aggregation into a population-scale scientific resource. The primary anticipated application is as an empirical verification layer for artificial intelligence: a continuously updated reference describing how human populations resolve comparable trade-offs, against which AI-generated decisions may be evaluated. Further applications span computational social science, behavioural economics, and public policy.

Background, Related Work & Gap

The scientific study of human decision-making spans multiple disciplines, including psychology, economics, sociology, political science, game theory, and artificial intelligence. Collectively, these fields have established many of the principles governing how individuals and groups make decisions, including bounded rationality, cognitive heuristics, strategic interaction, cooperation, reciprocity, social preferences, and ethical judgement. However, the methodologies employed within these disciplines have generally been developed to address specific research questions rather than to provide a continuously evolving empirical representation of collective decision-making.

Large-scale survey programmes have produced some of the most influential datasets describing human beliefs, values, and attitudes. Initiatives such as the **World Values Survey (WVS)**, **European Social Survey (ESS)**, **General Social Survey (GSS)**, and studies conducted by the **Pew Research Center** have established invaluable longitudinal resources for understanding cultural values, political attitudes, trust, religion, and social change. These programmes demonstrate the scientific value of systematic population-scale observation, but they primarily measure beliefs, attitudes, and opinions through periodically administered questionnaires rather than continuously observing structured decision-making across diverse contexts.

Experimental psychology and behavioural economics have similarly transformed our understanding of human decisions. Research by **Kahneman and Tversky** demonstrated the influence of cognitive heuristics and bounded rationality on decision-making, while subsequent work by **Thaler** and others established the importance of context, incentives, and framing effects. Behavioural ethics has further explored how organisational environments, incentives, identity, and social norms influence ethical decisions. These contributions have profoundly advanced decision science but have generally relied upon controlled experiments designed to test specific hypotheses rather than to construct continuously expanding observational datasets.

Game theory provides another important experimental tradition for studying decisions under strategic interaction. Classical paradigms including the Prisoner's Dilemma, Public Goods Game, Ultimatum Game, Dictator Game, Trust Game, Stag Hunt, and Common-Pool Resource experiments have generated foundational insights into cooperation, fairness, reciprocity, punishment, and collective action. More recently, online experimentation platforms have enabled these paradigms to be studied at substantially larger scales while maintaining experimental control.

Large-scale digital experiments have further demonstrated the scientific value of collecting structured decisions from globally distributed populations. The **Moral Machine** experiment, for example, gathered millions of responses to ethical trade-off scenarios across numerous countries, revealing substantial cultural variation in moral decision-making. Such work demonstrates that structured decision scenarios can generate valuable scientific insight when deployed at population scale. However, these studies have typically focused upon individual decision domains or specific experimental questions rather than operating as continuously evolving observational frameworks spanning multiple classes of decisions.

In parallel, computational social science has demonstrated the value of analysing large-scale digital datasets describing communication, mobility, commerce, and online activity. While these observational resources provide unprecedented insight into human activity, they primarily capture actions after decisions have already been made. Consequently, they provide only indirect evidence of the decision processes that produced those actions.

Despite these significant advances, an important methodological gap remains. Existing approaches generally fall into one of three categories: periodic surveys designed to measure beliefs and attitudes, controlled experiments designed to test specific hypotheses, or observational datasets describing naturally occurring human activity. None were designed as continuously operating observational infrastructures capable of collecting structured trade-off decisions across multiple domains, representing them within a common decision space, and analysing their evolution through time using a unified computational framework.

The framework presented in this work addresses this gap. It does not replace surveys, laboratory experiments, or observational datasets, nor does it propose a new theory of human decision-making. Instead, it introduces a continuously operating observational framework that aggregates structured decision observations into de-identified, population-level scientific representations. In doing so, it aims to complement existing methodologies by providing a reusable empirical resource for studying how populations make decisions across domains, how decision structures evolve through time, and how different classes of decisions relate to one another.

Framework Overview

Resonant Network or "Resonant" is a continuously operating instrument for observing collective human decision-making through the longitudinal collection of structured trade-off decisions. The platform and database described in this work are implemented and in production use, with a continuously growing base of registered users and recorded decision events; the framework is therefore presented as a functioning system rather than a design proposal. Current network statistics, including participation and accumulated decision events, are displayed in real time on the project website, through which this document is also accessible. Participants voluntarily engage with a diverse sequence of decision scenarios presented over time, each requiring an explicit choice between competing alternatives. The objective is not to infer psychological traits or to characterise individuals, but to observe how populations resolve trade-offs across a broad range of decision contexts and how those decision structures evolve through time.

Each completed scenario generates a **decision event** comprising the presented trade-off, the selected outcome, and associated de-identified contextual metadata. Contextual metadata is deliberately minimal, comprising only country-level location and observation time; no age, gender, or other demographic attributes are collected. This reflects a data-minimisation approach in which demographic resolution is deliberately traded for stronger privacy guarantees. Individually, these observations provide only limited scientific value. Collectively, however, they form a continuously expanding empirical representation of how populations make decisions under different conditions, across different domains, and over extended periods.

Unlike conventional survey instruments, Resonant is designed for continuous participation rather than single-session completion. Decision scenarios are intentionally distributed over time, allowing observations to be collected under changing personal, social, economic, and environmental conditions. This longitudinal design enables the framework to investigate not only how populations make decisions, but also how collective decision structures evolve in response to external events and changing contexts.

The primary measurement instrument consists of structured trade-off decisions spanning multiple domains, including everyday decisions, ethical dilemmas, strategic interactions, cognitive dilemmas, common-pool resource problems, organisational governance, healthcare, public policy, autonomous systems, financial decisions, and other domain-specific scenarios. Although these scenarios draw upon established survey and vignette methodologies, their purpose is not to measure isolated opinions or attitudes. Instead, they provide repeated observations of how participants resolve comparable decision trade-offs within a common framework.

The framework employs a predefined decision coordinate system: a base layer of six mutually orthogonal decision dimensions that provides a consistent mathematical basis for representing heterogeneous decision observations. Everyday trade-off decisions provide broad coverage of this coordinate system, establishing a common representation across participants and populations. More complex decision domains, including classical game-theoretic experiments, behavioural ethics scenarios, and multi-agent strategic interactions, form a second layer built on top of this base. Because such decisions do not reduce to the coordinate representation, they are analysed with non-traditional methods, such as multidimensional item response theory, topological data analysis, and the resulting structures are then compared against the underlying coordinate framework. This layered design enables relationships to be examined across fundamentally different classes of decisions while preserving the auditability of the base layer.

Participation simultaneously serves two independent purposes. At the individual level, participants progressively construct a personal decision profile derived from their own decision history. This profile remains under the participant's control and may be used for personal reflection, integrated with external AI systems, or applied within other decision-support contexts according to the participant's preferences. At the scientific level, only de-identified aggregate decision metadata contributes to the observational dataset. Scientific outputs take the form of population-level representations only; they are computed from individual participant records, which are de-identified and aggregated before any output is produced.

The framework consequently maintains a clear separation between **individual utility** and **scientific observation**. This dual-purpose architecture aligns participant incentives with scientific data collection, with the scientific contribution emerging from the aggregation of hundreds of thousands of independent decision events into population-level statistical structures.

Unlike traditional research datasets collected for a single study, the resulting observational resource is designed to expand continuously. As new decision events accumulate, the decision space becomes progressively richer, enabling increasingly detailed investigation of consensus, disagreement, decision dynamics, cross-domain relationships, and the evolution of collective decision structures at population scale.

Decision Representation and Computational Framework

Every completed trade-off within the Resonant framework generates a structured **decision event** that constitutes the fundamental unit of observation. A decision event is represented as

$$e = (s, r, t, c)$$

where

- **s** denotes the presented decision scenario,
- **r** denotes the participant's selected response,
- **t** denotes the observation timestamp, and
- **c** denotes the associated de-identified contextual metadata, comprising country-level location only.

Each decision event is independently observable and contains no interpretation regarding the participant's underlying cognition or intent. The framework measures decisions directly; any higher-order representations emerge only through the aggregation of many independent decision events.

Decision scenarios are organised around a predefined decision coordinate system whose base layer comprises six decision dimensions constructed to be mutually orthogonal. Rather than representing latent psychological constructs, these dimensions provide a common measurement basis that enables heterogeneous decision scenarios originating from different domains to be represented consistently within a shared space. At this foundational layer, each everyday trade-off scenario is mapped onto the six axes according to its predefined decision structure, so that completed responses yield coordinates within the resulting six-dimensional space. Importantly, the orthogonality of the base dimensions is a testable empirical property rather than a design assertion: because the framework is operational, it may be audited using conventional statistical techniques, such as correlation and factor analysis.

The choice of six dimensions reflects a deliberate balance between coverage and interpretability: the basis is intended as the largest practical set of axes capable of spanning the intended trade-off space while remaining individually interpretable and auditable. Whether six dimensions are sufficient, and whether the basis should be refined as evidence accumulates, is treated as an empirical question within the validation programme described below, rather than as a fixed design commitment.

Repeated observations across many decision events generate a decision coordinate for each participant at a given point in time. These participant-level coordinates are not themselves scientific outputs; they serve as intermediate computational representations from which aggregate population statistics are derived. All published or reported analyses are based on these aggregate representations, and identifiable participant trajectories are never presented as results.

For a population containing **N** participants, the aggregate distribution of decisions within the coordinate system may be represented by

$$\mu(t) = (1/N) \Sigma_{i=1}^{N} x_i(t)$$

where $\mathbf{x_i}$ denotes the decision coordinate associated with participant **i**.

The covariance structure of the observed population is given by

$$C(t) = (1/(N-1)) \Sigma_{i=1}^{N} (x_i(t) - \mu(t))(x_i(t) - \mu(t))^T$$

providing a quantitative description of the spread, diversity, and correlation of decisions across the observed population.

While everyday trade-off decisions primarily establish the underlying coordinate representation, more complex decision domains including game-theoretic interactions, ethical dilemmas, and other multi-agent decision problems introduce relationships that extend beyond independent coordinates. Unlike the base layer, these decision classes are characteristically multidimensional and collinear: a single decision loads simultaneously upon multiple dimensions and cannot be reduced to independent axis contributions. Such observations are therefore not forced into the six-dimensional basis. Instead, they are used either to construct higher-dimensional manifold representations of decision space or as inputs to MIRT and topological analysis, in which their relational structure is itself the object of study. These observations define an evolving **decision topology**, represented as

$$G = (V,E)$$

where **V** represents regions of similar aggregate decision structure and **E** represents statistically significant relationships between those regions inferred from the aggregate distribution of decision observations.

Importantly, this does not represent relationships between participants. Instead, it represents relationships between regions of collective decision space. As additional observations accumulate, the topology evolves continuously, revealing regions of consensus, persistent disagreement, decision clusters, and other emergent structures that arise directly from population-level observations.

Because Resonant operates continuously, every aggregate representation is inherently time-dependent. Let

$$G(t)$$

represent the topology of decision space at time **t**.

Structural evolution may therefore be expressed as

$$\Delta G = G(t_2) - G(t_1)$$

allowing investigation of decision drift, changes in consensus, responses to external events, and long-term evolution of collective decision structures.

Together, these quantities form the computational substrate of **Human Decision Intelligence (HDI)**: the continuously evolving mathematical description of collective human decisions, expressed as de-identified, population-level representations computed from individual decision events. Taken together, the centroid, covariance structure, density field, and evolving topology describe a single object of study: the geometry of collective human decision-making as it evolves through time. HDI is not conceived as a predictive model or theory of human cognition; its principal output classes are described in the Scientific Output section.

Verification of AI Decisions

Because decisions are inherently contextual and multidimensional, verification of an AI-generated decision is formulated as a retrieval and comparison problem rather than a lookup of any single scenario. Given a decision context presented to an AI system, the framework identifies the set of contextually proximate decision scenarios within the observational database, drawing upon the tens of thousands of decision contexts represented within the common decision space. The retrieval of this nearest contextual decision neighbourhood is deliberately specified only at this level of abstraction; no particular similarity metric or retrieval algorithm is prescribed by the framework. The population distribution of decisions over this proximate set, conditioned on observation time and on country as a societal proxy, provides the empirical reference against which the AI-generated decision is compared. Verification is therefore probabilistic rather than binary: an AI decision is assessed by its position within the observed population distribution for comparable contexts, rather than by agreement with any single majority outcome. Importantly, the framework is strictly descriptive: it measures what populations decided, not what is correct or ethical, and no normative judgement is encoded within the reference itself. A decision falling outside observed population norms is therefore not treated as an error in itself. Instead, divergence establishes a burden of explanation: the AI system under evaluation is expected to articulate the contextual considerations that motivated a decision outside the norms observed for comparable contexts. The framework thereby supports anomaly detection and explanation elicitation rather than adjudication, with the assessment of any explanation remaining a matter for the application domain.

Scientific Output

The Resonant framework transforms individual decision events into a hierarchy of aggregate scientific representations describing the structure, distribution, and evolution of collective human decision-making. Individual decision events are not themselves considered scientific outputs. Rather, they serve as de-identified observations that contribute to increasingly informative population-level representations as the observational dataset expands.

The framework generates several complementary classes of scientific output, each capturing a different aspect of collective decision space. Together, these outputs provide a continuously evolving empirical description of how populations make decisions across multiple domains without requiring analysis of identifiable individuals.

Population Decision Distributions

The most fundamental outputs are statistical distributions describing how decisions are made across populations. These distributions quantify the frequency of decisions within individual scenarios, decision domains, countries, and temporal intervals. They provide empirical measurements of consensus, disagreement, uncertainty, and variation, allowing collective decision structures to be studied quantitatively rather than anecdotally.

Unlike conventional survey summaries, these distributions evolve continuously as additional observations are collected, enabling longitudinal investigation of changing population decision patterns.

Decision Coordinate Representations

Basic decision events are mapped into a common six-dimensional coordinate system that provides a consistent mathematical representation across heterogeneous decision domains; complex decisions are handled in a separate analytical layer described below. Because all basic decisions share this reference space, aggregate populations, countries, cultures, and temporal cross-sections can be represented as distributions within the same coordinate space, enabling quantitative comparison between groups that have traditionally been studied independently.
The coordinate system is intended as a measurement framework rather than a taxonomy of human cognition. Its purpose is to provide a stable, auditable reference space within which basic decisions can be represented consistently.

Decision Topology

As observations accumulate, relationships emerge between different regions of decision space. Complex decision domains, including game-theoretic experiments, behavioural ethics scenarios, and multi-agent strategic interactions, are mapped directly into this layer rather than into the coordinate system, and are analysed using non-traditional methods.

Together, these produce an evolving decision topology describing the structural organisation of collective decisions: regions of consensus, persistent disagreement, transition pathways, clusters of related decision patterns, and other higher-order structures. Because the topological layer and the coordinate system share the same underlying decision events, structures identified here can be compared against the auditable base layer, allowing relationships to be examined across fundamentally different classes of decisions.

Decision Dynamics

Because Resonant operates as a continuously evolving observatory rather than a periodic survey, changes in collective decisions may be observed directly through time.

Longitudinal observation enables investigation of:

- decision drift,
- consensus formation,
- fragmentation,
- responses to significant external events,
- structural stability,
- and the emergence of previously unobserved decision patterns.

This provides a dynamic representation of collective decision-making that cannot be obtained from isolated surveys or one-off experiments.

Cross-Domain Decision Relationships

A distinguishing feature of the framework is that every decision domain is represented within the same observational system.

Consequently, statistical relationships may be explored between decisions originating from fundamentally different contexts, including everyday decisions, strategic interactions, ethical dilemmas, financial decision-making, and specialised domains.

This enables investigation of whether common decision structures exist across domains that have traditionally been studied independently.

Emergent Decision Archetypes

Rather than assigning predefined psychological categories or personality types, Resonant allows recurring patterns of decisions to emerge empirically from aggregate observations.

Decision archetypes are recurring aggregate patterns of decision-making, identified as characteristic distributions within the coordinate space that appear consistently across multiple decision domains. Because archetypes are re-derived from continuously updated population-level representations, the classification evolves with the underlying population rather than being fixed in advance.

Personal Decision Profiles

While all scientific outputs consist of aggregate, de-identified decision metadata, each participant simultaneously develops a personal longitudinal decision profile based on their own decision history.

Unlike the population-level scientific outputs described above, these personal profiles are intended for individual use. They enable participants to explore the consistency of their own decisions through time, compare their decisions with aggregate population structures, and, where they choose, integrate their personal decision profile with external AI systems or other decision-support applications. Participants are therefore not merely contributors of observations: each accumulates a durable, participant-owned asset whose value grows with continued participation.

This dual architecture separates **individual utility** from **scientific observation**. Personal decision profiles remain participant-centric, while the scientific observatory reports only aggregate decision structures.

Human Decision Intelligence

Taken together, these scientific outputs constitute the empirical foundation of **Human Decision Intelligence**.

Rather than producing isolated survey statistics or experimental results, the framework generates a continuously expanding representation of collective human decision-making composed of population decision distributions, decision coordinate representations, evolving decision topologies, temporal decision dynamics, cross-domain decision relationships, and emergent decision archetypes. These outputs differ in auditability: population distributions and coordinate representations remain traceable to the six interpretable axes, whereas derived representations such as decision-space embeddings and evolving topologies rely on higher-dimensional or learned structures that cannot be traced back to individual axes.

Importantly, every scientific output consists exclusively of aggregate, de-identified decision metadata; individual-level records serve only as intermediate inputs to these aggregates. Individual participants are not the object of scientific analysis. Instead, each decision contributes, in de-identified form, to larger statistical structures describing populations rather than individuals.

As the observational dataset expands, these representations become progressively richer, enabling increasingly detailed investigation of how populations make decisions, how decision structures evolve, and how collective decision-making changes across contexts, cultures, and time.

Limitations and Future Work

This framework is a measurement infrastructure, not a theory of human decision-making. Several methodological questions therefore remain open and will need continued attention as the framework develops.

First, participation is voluntary and therefore subject to the limitations common to large-scale digital platforms. The observed population should not be assumed, at least initially, to constitute a statistically representative sample of any national or global population. Consequently, interpretation of aggregate decision structures should account for potential participation biases arising from geography, culture, language, technology adoption, and demographic composition. As the observational dataset expands, established statistical weighting and sampling methodologies may be incorporated where appropriate to improve representativeness for specific research applications. Because demographic attributes are not collected, such participation biases cannot be directly quantified within the dataset itself. For the same reason, observed longitudinal changes in aggregate decision structures may reflect changes in the composition of the participant base as well as genuine change within a stable population; panel-style analysis of continuing participants provides one means of distinguishing these effects.

Second, the framework measures decisions made within structured trade-off scenarios rather than decisions made under real-world consequences. While structured decision scenarios have a long history within survey research, behavioural economics, experimental psychology, and game theory, hypothetical decisions cannot be assumed to reproduce all aspects of real-world decision-making. The framework instead provides a continuously operating resource for measuring how populations resolve comparable decision trade-offs under controlled conditions. Future work should investigate the relationship between observed decision structures and real-world decision outcomes across different application domains.

Third, the dual-purpose architecture introduces a distinctive incentive consideration. Because participants progressively construct a personal decision profile that they may review, compare against aggregate population structures, and integrate with external decision-support tools, responses may be influenced by self-presentation and social desirability effects, with participants answering as they would wish to decide rather than as they would decide in practice. Such effects are common to self-report instruments but may be amplified where the instrument itself rewards self-curation. Potential mitigations include consistency checks across structurally equivalent scenarios, and comparison of decision distributions across stages of participation; these approaches form part of the framework's ongoing validation programme.

Fourth, the quality of the resulting scientific resource depends upon the design and validation of the underlying decision scenarios. Decision wording, cultural interpretation, linguistic equivalence, and measurement invariance across populations all influence the comparability of observations. Continuous validation, refinement, and expansion of the scenario corpus therefore constitute an integral part of the framework's long-term development.

The decision coordinate system similarly represents a measurement framework rather than a fixed scientific model. The number, definition, and weighting of decision vectors may evolve as additional empirical evidence becomes available, and alternative computational representations may prove advantageous for particular classes of decisions or analytical objectives. Likewise, the computational methods used to estimate decision embeddings, identify decision topologies, and detect emergent structures may continue to develop as the dataset increases in scale and diversity.

Validation Programme

The scientific standing of the framework depends on demonstrating that it measures something real, and this is addressed through an explicit validation programme rather than through design assertions alone. The programme comprises empirical testing of axis independence through correlation and factor analysis of the accumulated corpus; test-retest reliability of individual scenarios; consistency of responses across structurally equivalent trade-offs; temporal stability of aggregate distributions in the absence of significant external events; and convergence with established behavioural and survey measures where decision domains overlap. Several of these streams were introduced above in the context of specific limitations; they are consolidated here because together they determine whether the coordinate system, the topological layer, and the derived representations can be trusted as measurements.

Finally, while all scientific outputs are aggregate, de-identified metadata, the operational platform necessarily retains participant decision histories in order to generate personal longitudinal decision profiles. Future work will examine privacy-preserving computation, governance, and participant control over data sharing, with the aim of serving individual participants without compromising population-scale observation.

These considerations do not diminish the underlying objective of the framework. Rather, they reflect the normal methodological evolution of any new scientific instrument. As additional observations accumulate, computational methods mature, and validation studies are performed, the framework is expected to evolve alongside the questions it is designed to support.

Impact & Applications

A continuously operating observatory of structured human decisions has the potential to provide a new empirical resource for scientific research, public policy, and intelligent decision-support systems. By transforming large numbers of individual decision events into aggregate statistical representations, the framework enables investigation of collective decision-making at spatial, temporal, and contextual scales that have previously been difficult to observe using conventional methodologies.

The primary intended application lies in artificial intelligence. Contemporary AI systems are trained primarily on language, images, software, and records of human activity, while comparatively little structured information exists describing how large populations resolve comparable decision trade-offs across multiple domains. Population-scale decision metadata therefore provides a complementary empirical resource: the verification layer described earlier, against which AI decision processes may be evaluated, supporting alignment research and the development of AI systems capable of reasoning more effectively about human decisions. For this class of application, population-level decision structures are the relevant reference, and demographic resolution is not required.

Beyond this primary application, the framework provides the social sciences with a common infrastructure for studying how populations resolve trade-offs across diverse domains. Rather than examining individual decision scenarios in isolation, researchers may investigate how decision structures evolve through time, how consensus emerges or fragments, how decisions vary across countries or cultures, and how external events influence collective decision-making. Because basic decisions share a common coordinate system and complex decisions occupy a linked topological layer within the same observational framework, relationships between previously independent areas of decision research may also be explored within a unified analytical setting.

The framework also provides a foundation for computational approaches to decision science. Aggregate decision structures may be analysed using statistical learning, network science, topology, dimensionality reduction, and other computational techniques to investigate the geometry of decision space, identify persistent decision structures, quantify decision diversity, and explore relationships between different classes of decisions.

Beyond scientific research, the framework enables participants to derive direct personal value from their own decision histories. This dual-purpose architecture aligns participant benefit with the continuous growth of the scientific observatory.

More broadly, the framework may support evidence-informed decision-making across public policy, organisational governance, healthcare, education, economics, and other domains in which understanding collective decisions is important. It provides a continuously expanding empirical representation of how populations resolve structured trade-offs through time.

Conclusion

This work has described a continuously operating observational framework for studying collective human decision-making through the longitudinal collection of structured trade-off decisions. Rather than replacing existing surveys, laboratory experiments, or observational datasets, the framework is intended to complement these methodologies by providing a common empirical representation through which diverse classes of decisions may be observed, compared, and analysed over time.

The principal contribution of the framework is not a new theory of human cognition or behaviour. Instead, it introduces a general-purpose infrastructure that transforms individual decision events into de-identified, population-level scientific representations, including statistical distributions, decision coordinate representations, evolving decision topologies, and longitudinal measures of collective decision dynamics. These representations enable the study of populations rather than individuals while preserving participant privacy by publishing only aggregate decision metadata.

A distinguishing characteristic of the framework is the alignment between individual utility and scientific observation: individual benefit and scientific value develop together rather than competing for participant engagement.

The framework is intended to evolve alongside the dataset that it generates. As participation expands, decision scenarios mature, computational methods improve, and additional validation studies are performed, the resulting scientific resource may enable increasingly detailed investigation of how populations resolve trade-offs across cultures, domains, and time. In particular, it is intended to serve as a continuously updated empirical reference against which the decisions of artificial intelligence systems may be verified, providing a population-grounded benchmark for evaluating whether AI-generated decisions remain consistent with observed human decision structures. Future work will determine the extent to which these observations contribute to decision science, computational social science, artificial intelligence, public policy, and other disciplines concerned with understanding collective human decisions.

Ultimately, the contribution of this work lies in demonstrating that structured human decisions can be observed continuously, represented within a common computational framework, and studied as an evolving population-scale phenomenon. If successful, such an observatory may provide a reusable empirical foundation for future research into collective decision-making while simultaneously delivering practical value to the individuals whose decisions make that observation possible.

References

Decision Science and Behavioural Economics

1. Kahneman, D., & Tversky, A. (1979). *Prospect Theory: An Analysis of Decision under Risk*. *Econometrica*, 47(2), 263–291.
2. Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under Uncertainty: Heuristics and Biases*. Cambridge University Press.
3. Simon, H. A. (1955). *A Behavioral Model of Rational Choice*. *Quarterly Journal of Economics*, 69(1), 99–118.
4. Simon, H. A. (1957). *Models of Man*. Wiley.
5. Thaler, R. H., & Sunstein, C. R. (2008). *Nudge*. Yale University Press.

Behavioural Ethics

6. Bazerman, M. H., & Tenbrunsel, A. E. (2011). Blind Spots: Why We Fail to Do What's Right and What to Do about It. Princeton University Press.
7. Kaptein, M. (Ed.). (2015). Behavioral Ethics in Organizations. Edward Elgar.

Game Theory

8. von Neumann, J., & Morgenstern, O. (1944). Theory of Games and Economic Behavior. Princeton University Press.
9. Axelrod, R. (1984). The Evolution of Cooperation. Basic Books.
10. Osborne, M. J. (2004). An Introduction to Game Theory. Oxford University Press.

Survey Research and Population Studies

11. Haerpfer, C., et al. (2022). World Values Survey: Round Seven (2017–2022).
12. European Social Survey. Methodology and Scientific Documentation.
13. Smith, T. W., Marsden, P., Hout, M., & Kim, J. General Social Survey.
14. Pew Research Center. Methodological Reports and Survey Documentation.
15. Aguinis, H., & Bradley, K. J. (2014). Best Practice Recommendations for Designing and Implementing Experimental Vignette Methodology Studies. Organizational Research Methods, 17(4), 351–371.
16. Tourangeau, R., & Yan, T. (2007). Sensitive Questions in Surveys. Psychological Bulletin, 133(5), 859–883.

Computational Social Science

17. Lazer, D., Pentland, A., Adamic, L., et al. (2009). Computational Social Science. Science, 323, 721–723.
18. Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., Bonnefon, J.-F., & Rahwan, I. (2018). The Moral Machine Experiment. Nature, 563, 59–64.

Network Science

19. Newman, M. E. J. (2010). Networks: An Introduction. Oxford University Press.
20. Barabási, A.-L. (2016). Network Science. Cambridge University Press.
21. Kazemi, S. M., Goel, R., Jain, K., Kobyzev, I., Sethi, A., Forsyth, P., & Poupart, P. (2020). Representation Learning for Dynamic Graphs: A Survey. Journal of Machine Learning Research, 21(70), 1–73.

Topological Data Analysis

22. Carlsson, G. (2009). Topology and Data. Bulletin of the American Mathematical Society, 46(2), 255–308.
23. Giusti, C., Pastalkova, E., Curto, C., & Itskov, V. (2015). Clique Topology Reveals Intrinsic Geometric Structure in Neural Correlations. Proceedings of the National Academy of Sciences, 112(44), 13455–13460.

Dimensionality Reduction and Representation Learning

24. Tenenbaum, J. B., de Silva, V., & Langford, J. C. (2000). A Global Geometric Framework for Nonlinear Dimensionality Reduction. Science, 290, 2319–2323.
25. Belkin, M., & Niyogi, P. (2003). Laplacian Eigenmaps for Dimensionality Reduction and Data Representation. Neural Computation, 15(6), 1373–1396.
26. Bengio, Y., Courville, A., & Vincent, P. (2013). Representation Learning: A Review and New Perspectives. IEEE Transactions on Pattern Analysis and Machine Intelligence, 35(8), 1798–1828.

Artificial Intelligence

27. Russell, S. (2019). Human Compatible. Viking.
28. Miller, T. (2019). Explanation in Artificial Intelligence: Insights from the Social Sciences. Artificial Intelligence, 267, 1–38.

29. Christiano, P. F., Leike, J., Brown, T. B., Martic, M., Legg, S., & Amodei, D. (2017). Deep Reinforcement Learning from Human Preferences. In Advances in Neural Information Processing Systems, 30.
30. Hendrycks, D., Burns, C., Basart, S., Critch, A., Li, J., Song, D., & Steinhardt, J. (2021). Aligning AI with Shared Human Values. In International Conference on Learning Representations.

Privacy

31. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data (General Data Protection Regulation).
32. Dwork, C. (2006). Differential Privacy. In Automata, Languages and Programming (ICALP 2006), Lecture Notes in Computer Science, 4052, 1–12. Springer.
33. Sweeney, L. (2002). k-Anonymity: A Model for Protecting Privacy. International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, 10(5), 557–570.