



D3.4: Report on virtual worlds environments exploratory study

Revision: v.1.0

Work package	WP 3
Task	4
Due date	31/10/2025
Submission date	31/10/2025
Deliverable lead	The Lisbon Council
Version	1.0
Authors	Francesco Mureddu (LC)
Reviewers	Francesco Panella (MARTEL); Shani Tiran (IDSA)
Abstract	This deliverable reports an exploratory, controlled experiment run inside two mainstream social-VR platforms, Horizon Worlds and VRChat, to test how visible governance safeguards (“externally signalled rules”) and digital-identity posture shape user trust, perceived openness, immersion, and comfort. The 2x2 design crossed Rule Framework (platform-only norms vs. externally signalled rules) with Digital Identity (pseudonymous/anonymous vs. visible verified cue); sessions used Meta Quest headsets and a common onboarding–exploration–survey protocol (N=101, balanced across cells). The primary outcomes were trust/privacy and openness; secondary outcomes were spatial presence and Simulator Sickness (SSQ). Three signals stand out. First, a robust onboarding effect: trust/privacy increased from pre- to post-exposure across the full sample ($t(100) = -3.70$, $p < .001$; $d \approx 0.53$). Second, visible safeguards improved outcomes without “closing” the world: externally signalled rules raised trust/privacy ($F(1,96)=5.02$, $p=.027$) and markedly reduced simulator sickness ($F(1,96)=10.70$, $p=.001$), with no significant decline in perceived openness. Third, making identity more visible did not add trust or openness but reduced spatial presence ($F(1,96)=4.85$, $p=.030$), indicating an experiential cost to real-name/verified cues in short interactions. Correlations reinforced comfort as a hinge variable: higher trust/privacy aligned with lower sickness and higher openness. Implications for policy and uptake follow directly: prioritise visible safeguards (trust labels and pre-entry rule/redress summaries), a safety-by-design UX baseline, privacy-preserving credentials (selective-disclosure VCs) over identity exposure, and

www.open-verse.eu



Grant Agreement No.: 101135701
Call: HORIZON-CL4-2023-HUMAN-01-CNECT

Topic: HORIZON-CL4-2023-HUMAN-01-23
Type of action: HORIZON-CSA

	auditable safety manifests, validated and scaled through pan-EU sandboxes and procurement.
Keywords	Virtual worlds (social VR); Governance signalling; Digital identity (privacy-preserving credentials); Trust & privacy; Simulator sickness.

Document Revision History

Version	Date	Description of change	List of contributor(s)
V0.1	15/09/2025	ToC and abstract	Francesco Mureddu (LC)
V0.9	13/10/2025	First draft	Francesco Mureddu (LC)
V0.9	13/10/2025	Internal review for quality assurance	Francesco Mureddu (LC)
V1.0	31/10/2025	Final version for submission	Francesco Mureddu (LC)

DISCLAIMER



**Funded by
the European Union**

Project funded by



Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Federal Department of Economic Affairs,
Education and Research EAER
State Secretariat for Education,
Research and Innovation SERI

Funded by the European Union (OPENVERSE, 101135701). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

This work has received funding from the Swiss State Secretariat for Education, Research and Innovation (SERI).

COPYRIGHT NOTICE

© 2023 - 2026 OPENVERSE

Project funded by the European Commission in the Horizon Europe Programme		
Nature of the deliverable:	R	
Dissemination Level		
PU	Public, fully open, e.g. web (Deliverables flagged as public will be automatically published in CORDIS project's page)	✓
SEN	Sensitive, limited under the conditions of the Grant Agreement	
Classified R-UE/ EU-R	EU RESTRICTED under the Commission Decision No2015/ 444	
Classified C-UE/ EU-C	EU CONFIDENTIAL under the Commission Decision No2015/ 444	
Classified S-UE/ EU-S	EU SECRET under the Commission Decision No2015/ 444	

* R: Document, report (excluding the periodic and final reports)



**Funded by
the European Union**

Project funded by

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Federal Department of Economic Affairs,
Education and Research EAER
State Secretariat for Education,
Research and Innovation SERI

DEM: Demonstrator, pilot, prototype, plan designs

DEC: Websites, patents filing, press & media actions, videos, etc.

DATA: Data sets, microdata, etc.

DMP: Data management plan

ETHICS: Deliverables related to ethics issues.

SECURITY: Deliverables related to security issues

OTHER: Software, technical diagram, algorithms, models, etc.

EXECUTIVE SUMMARY

OPENVERSE set out to answer a simple, public-interest question: can virtual worlds feel safer and more trustworthy without becoming less open and creative? To move beyond opinion, we ran a **controlled experiment inside two mainstream social VR platforms** (i.e. Horizon Worlds and VRChat), so the effects of governance signalling and digital-identity posture could be measured directly on real users. The design crossed **two factors, Rule Framework and Digital Identity**, to create four conditions: platform-only norms vs. externally signalled rules, and pseudonymous/anonymous use vs. a visible verified identity cue. **Two orthogonal stressors, visual fidelity and interaction complexity**, were varied where feasible to check that results were not artefacts of graphics style or task demands. Participants used standard Meta Quest headsets and followed the same onboarding, exploration, and survey sequence across conditions.

The analytic sample comprised **101 participants** with balanced randomisation across the four cells (approximately twenty-five per cell) and **56% female**. All participants completed pre- and post-experience questionnaires, yielding a clean dataset with no material missingness after standard data checks. Scale reliability was **good** for Openness ($\alpha \approx .77$), **acceptable** for Trust/Privacy ($\alpha \approx .63$), and **excellent** for Simulator Sickness ($\alpha \approx .91$); Presence ($\alpha \approx .24$) and baseline privacy ($\alpha \approx -.09$ after removing invariant items) were flagged as developmental instruments, guiding future refinement rather than undermining the headline effects.

Three results stand out with statistical clarity. First, there is a robust **onboarding effect**: a paired analysis shows **significant pre→post gains in Trust/Privacy** across the full sample ($t(100) = -3.70, p < .001$, **Cohen's d = 0.53**). A short, structured introduction to VR through clear instruction and a brief acclimatisation, makes people feel safer and more confident, regardless of platform. Second, **externally signalled rules** measurably improve outcomes. Compared with platform-only settings, externally signalled environments registered a **significant main effect on Trust/Privacy** ($F(1,96) = 5.02, p = .027$) and a **substantial reduction in Simulator Sickness** ($F(1,96) = 10.70, p = .001$), while **perceived Openness did not decline** (no significant differences). In plain terms, visible safeguards lifted trust and cut discomfort, often by the order of **forty to fifty percent** on sickness indices, without “closing” the world. Third, **making identity more visible** did **not** add trust or openness and **reduced spatial Presence** ($F(1,96) = 4.85, p = .030$), indicating a tangible user-experience cost to real-name/verified cues with no compensating trust benefit in short interactions.

The correlational structure explains why these patterns matter for everyday users. **Trust/Privacy correlated positively with Openness** ($r = .61, p < .001$) and **negatively with Simulator Sickness** ($r = -.43, p < .001$): when people feel safer and better treated, the world also feels more open and less physically taxing. **Presence correlated positively with sickness** ($r = .32, p < .01$) and **negatively with trust/privacy** ($r = -.26, p < .01$), a reminder that pushing immersion without commensurate comfort features can erode the sense of safety. Importantly, **VR literacy** (i.e. prior exposure to VR and gaming) did not drive these outcomes, which strengthens the case that effects come from the policy-relevant levers, not from who happened to be in each group.

For policymakers and the public, the message is straightforward. **Visible safeguards work**. If rules and redress are clearly signposted before entry and reinforced inside the world, trust rises and discomfort falls, while creativity and perceived openness remain intact. **Identity exposure does not**. Forcing real-name visibility doesn't deliver the hoped-for trust gain and it compromises immersion; accountability should be provided through **privacy-preserving credentials** (proving age, eligibility or uniqueness without revealing personal data) and **clear, explainable redress**, not through exposing identities. Because the largest experiential gain is reduced sickness, comfort features are not “nice-to-have”, they

are an inclusion lever that keeps more people, including those sensitive to motion or social risk, engaged.

The study also delivers a replicable method and public-facing artefacts. We document a transparent analysis plan (paired pre → post estimates; ANCOVA with covariates; mixed-effects checks), item-level data cleaning (removal of zero-variance items; reverse-key recoding), and balanced randomisation, producing effect sizes and confidence intervals that can be reused in future evaluations. This means the findings are not a one-off: they can be folded into standards, procurement, and sandboxes so that “what works” becomes “what’s required.”

Translating evidence into recommendations and action is now a matter of packaging what users respond to. A lightweight **trust label** that presents rules, user rights, and redress **before entry**; a small **safety-by-design baseline** (panic/exit, blocking and proximity controls, harassment filters, sickness-aware locomotion, session-time prompts); a **machine-readable safety manifest** with explainable moderation/AI logs; and **verifiable-credential flows** for accountability without identity exposure together reflect the data and can be tested under a common sandbox charter. Publishing certification status, conformance results, and pilot logs to a public observatory keeps the evidence loop open and builds legitimacy.

In sum, the experiment shows that **clear rules beat exposed identities** as a path to trust. By making safeguards visible, engineering comfort as a baseline, keeping accountability privacy-preserving, and documenting protections in ways people can verify, Europe can grow virtual worlds that are safer and more welcoming **without** sacrificing openness and imagination.

TABLE OF CONTENT

Disclaimer	2
Copyright notice	2
EXECUTIVE SUMMARY	4
LIST OF FIGURES AND LIST OF TABLES	8
ABBREVIATIONS.....	9
INTRODUCTION.....	11
1 BASELINE AND OBJECTIVES	14
1.1 Key variables	14
1.2 Area of benchmarking (technological, ethical, legal, equality	15
1.3 Selected Virtual World.....	16
2 METHODOLOGY OF THE EXPERIMENT	18
2.1 Participants	18
2.2 Materials and apparatus	18
2.3 Procedure.....	20
2.4 Measures.....	21
2.5 Statistical measures	21
3 RESULTS OF THE EXPERIMENTS.....	24
3.1 Scale Reliability and Data Cleaning.....	24
3.2 Pre-Post Changes in Trust and Privacy Perceptions and effects of manipulations ..	25
3.3 Correlational Patterns Among Outcome Variables.....	26
4 BENCHMARK CONCLUSIONS.....	28
4.1 How the experimental results validate (and nuance) the foresight track in D2.3	28
4.2 Benchmark against “Trust Tools for Immersive Technologies” (WP3).....	28
4.3 Implications for OPENVERSE roadmapping (T2.5)	29
5 DISCUSSION	31
5.1 Further research.....	31
5.2 Use of results of the experiment in future OPENVERSE activities.....	32
5.2.1 From “wild west” spaces to gently structured environments	32
5.2.2 Identity visibility versus pseudonymous accountability	32
5.2.3 Openness and interoperability, tempered by safety baselines	33
5.2.4 Synthesis and implications for T2.4 facilitation	33
5.2.5 Closing note	34
5.3 Policy recommendations.....	34
6 CONCLUSIONS.....	38

REFERENCES.....	39
ANNEX I: THE TWO VIRTUAL WORDS USED IN THE EXPERIMENT	40
ANNEX II: THE HEADSET USED IN THE EXPERIMENTS	42
ANNEX III: EXPERIMENTS' QUESTIONNAIRES AND EVALUATION SCALES.....	43
ANNEX IV: ETHICS ASSESSMENT BY ETHICS ADVISOR	1
6.1 Ethics Advisor Assessment Report – Deliverable D3.4	1
6.1.1 Overview	1
6.1.2 Summary of Ethical Compliance	1
6.1.3 Sensible Points and Mitigation Measures	1
6.1.4 Overall Ethical Acceptability.....	2
6.1.5 Recommendations for Future Actions	2
6.1.6 Conclusion	2
ANNEX V: DPIA SUMMARY	3
6.2 Data Protection Impact Assessment (DPIA) Summary – Deliverable D3.4	3
6.2.1 Purpose of the DPIA	3
6.2.2 Description of the Processing	3
6.2.3 Risk Assessment	3
6.2.4 Data Flow and Storage	4
6.2.5 Data Retention and Deletion	4
6.2.6 Consultation and Approval	4
6.2.7 Conclusion	4

LIST OF FIGURES AND LIST OF TABLES

FIGURE 1: THE TWO DIMENSIONS OF THE EXPERIMENTS AND THE 4 QUADRANTS.....	19
FIGURE 2: 3D SPACE FOR TWO TECHNICAL FACTORS AND THE TWO MAIN DIMENSIONS OF THE ANALYSIS.....	20
FIGURE 3 SCREENSHOT FROM THE EXPERIMENT	41
TABLE 1: SCALE RELIABILITY AND ITEM RETENTION	24
TABLE 2: INTERCORRELATIONS AMONG PRIMARY OUTCOME VARIABLES	26

ABBREVIATIONS

AI	Artificial Intelligence
ANOVA	Analysis of Variance
ANCOVA	Analysis of Covariance
CORDIS	Community Research and Development Information Service
CSA	Coordination and Support Action
DATA	Data sets / microdata (deliverable type)
DEC	Websites, patents filing, press & media actions, videos, etc. (deliverable type)
DEM	Demonstrator, pilot, prototype, plan designs (deliverable type)
DID	Decentralized Identifier
DMP	Data Management Plan (deliverable type)
eIDAS	Electronic Identification, Authentication and Services
EC	European Commission
EIB	European Investment Bank
EIF	European Investment Fund
EU	European Union
GDPR	General Data Protection Regulation
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IPQ	Igroup Presence Questionnaire
ISO	International Organization for Standardization
ITU	International Telecommunication Union
KPI	Key Performance Indicator
MoU	Memorandum of Understanding
MS	Member States
SDK	Software Development Kit
SDO	Standards Developing Organization
SME	Small and Medium-sized Enterprise
SSQ	Simulator Sickness Questionnaire
UX	User Experience

VC	Verifiable Credential
VR	Virtual Reality
W3C	World Wide Web Consortium
XR	Extended Reality

INTRODUCTION

OPENVERSE project advances a European vision of virtual worlds that are open, interoperable, and human-centric by combining foresight, empirical testing, and policy translation. Within this framework, Task T3.2 was mandated to conduct an exploratory study of existing virtual-world environments and to generate user-level evidence on behaviour, trust, privacy, and digital identity. To that end, we designed and ran a small-scale experiment in two mainstream social VR platforms (Meta's Horizon Worlds and VRChat) under systematically varied governance and identity conditions. The experiment was never intended as a market comparison; rather, it operationalised contrasting governance cues and identity postures so that their effects on user experience could be measured transparently and mapped back to the OPENVERSE foresight axes. The resulting evidence informs roadmapping and policy options under T2.5, thereby closing the loop from anticipatory analysis and foresight made in T2.4 to implementation guidance for the OPENVERSE WP5 and WP6.

The study is anchored in a clear question that echoes current European strategy for Web 4.0 and virtual worlds: do visible, externally signalled rules meaningfully improve perceived safety and comfort without narrowing openness, and what role, if any, does overt identity verification play in building trust? By isolating two contextual levers, a rule framework (platform-only norms versus externally signalled oversight) and a digital identity policy (pseudonymous/anonymous use versus a verified/real-name cue), the experiment translates foresight hypotheses into testable contrasts with interpretable effect sizes. A pair of orthogonal technical stressors visual fidelity and interaction complexity were added to probe ecological sensitivity while keeping hardware, locomotion, and session structure constant. The outcome set combines policy-salient constructs (trust/privacy; perceived openness) with experiential safety and quality (presence/immersion; simulator sickness), allowing us to relate user experience directly to governance signals and identity posture.

Methodologically, this work packages a 2x2 factorial design into real platforms with controlled cue design and exposure. Participants completed baseline questionnaires, a short acclimatisation, a guided exploration in their assigned condition, and a post-experience battery. The protocol emphasised internal validity through balanced randomisation, item-level data cleaning, and reliability checks, while preserving external face validity by using recognisable environments and platform-native affordances. The analysis plan combines pre/post change estimation with between-group inference (ANOVA/ANCOVA and confirmatory mixed-effects models), reports effect sizes and confidence intervals, and inspects assumptions and robustness. This produces results that are both statistically defensible and decision-useful for downstream tasks.

Substantively, the emerging signals are straightforward and actionable. Externally signalled rule frameworks lift perceived trust/privacy and markedly reduce simulator sickness without depressing perceived openness, suggesting that visible safeguards and legible redress mechanisms are high-yield levers that do not “close” creative ecosystems. By contrast, visible identity verification does not add trust or openness and tends to depress presence, indicating that accountability is better achieved via privacy-preserving credentials and due-process-ready redress than through blanket real-name display. Correlational patterns reinforce comfort as a hinge variable: as discomfort falls, trust rises; as immersion is pushed to its limits, discomfort can increase. These findings align with OPENVERSE's foresight narrative on “sovereign openness” and provide concrete handles for policy tools such as trust labels, selective-disclosure credentials, interoperable safety UX baselines, and auditability artefacts.

This report, therefore, has a dual mission. First, it documents the objectives, materials, procedures, measures and statistical strategy of the T3.2 experiment, presenting the results in a form suitable for academic scrutiny and replication. Second, it reads those results against the project's foresight scenarios and the evolving European policy canvas (T2.4), extracting implications for the project

workplan and beyond: codifying visible safeguards into a recognisable label, specifying privacy-preserving identity proofs compatible with eIDAS-aligned wallets, defining a minimum safety-by-design user-experience baseline, institutionalising machine-readable safety manifests and explainable moderation trails, and organising cross-border sandboxes so that supervision and testing travel together. In doing so, the document offers a coherent path from scenario to specification, from specification to certification, and from certification to uptake, keeping user experience at the centre while equipping T2.5, WP6 and WP5 with evidence they can act on immediately.

This report received prior ethical clearance from the OPENVERSE Ethics Advisor in line with D7.7 and D7.3 requirements (see Annex IV).

The document is organized as follows:

Chapter 1: Baseline and Objectives.. This chapter sets the evidentiary baseline and the concrete objectives that anchor the experiment-to-policy pipeline. It explains why governance signalling and digital identity were selected as the pivotal contextual levers, how they were operationalised in two mainstream social-VR platforms under tightly controlled technical conditions, and which outcome constructs translate foresight questions into testable hypotheses and interpretable effect sizes. It also defines the benchmarking perimeter across technological, ethical, legal and equality domains so that results can travel into standards, procurement and roadmapping. World selection (Horizon Worlds; VRChat) and cue-embedding choices are documented with pointers to Annex I.

Chapter 2: Methodology of the Experiment. This chapter details the 2x2 factorial design such as Rule Framework (platform-only norms vs externally signalled rules) × Digital Identity (pseudonymous/anonymous vs visible verified cue), and the exploratory technical stressors (visual fidelity; interaction complexity). It specifies participant recruitment and randomisation (N=101), headsets and platform builds, onboarding and exposure protocol, instruments (primary: trust/privacy, openness; secondary: presence/IPQ, simulator sickness/SSQ), data-cleaning rules (zero-variance removal; reverse-key recoding), and the statistical strategy (paired pre/post; ANCOVA with covariates; mixed-effects checks; effect sizes; robustness diagnostics). Together, these sections show precisely how the governance and identity manipulations were operationalised and how inferences were derived.

Chapter 3: Results of the Experiments. Findings are presented in four blocks: (i) scale reliability and data cleaning, (ii) pre → post changes (onboarding gains in trust/privacy), (iii) main and interaction effects of Rule Framework and Digital Identity on post-test endpoints (trust/privacy, openness, presence, SSQ), and (iv) correlational structure among outcomes. In brief: externally signalled rules significantly raised trust/privacy and markedly reduced simulator sickness without depressing perceived openness; visible identity verification did not add trust/openness and lowered spatial presence; VR literacy did not drive effects.

Chapter 4: Benchmark Conclusions. This chapter synthesises the empirical signals against OPENVERSE's foresight logic ("sovereign openness"). It explains why visible safeguards (rule legibility, redress) increased trust and comfort without closing worlds, while identity exposure failed to add trust and reduced immersion. It clarifies limitations and generalisation boundaries to guide cautious but useful policy reading and draws explicit links to roadmapping and standards pathways.

Chapter 5: Discussion. This chapter bridges evidence and execution. It first outlines a practical research agenda to deepen confidence in the signals observed: enlarge and diversify samples; refine instruments where reliability was borderline (especially presence and baseline privacy scales); test persistence through extended and repeated exposures; and broaden platform coverage and technical stressors to probe generalisability. In parallel, it translates today's results into near-term actions across

the project: informing scenario design and experimentation logic; shaping the roadmap so that safeguards are treated as product features rather than afterthoughts; and guiding upcoming work packages to privilege privacy-preserving accountability over identity exposure, and gentle governance cues over heavy-handed controls. The chapter then concentrates the policy and technical levers that flow directly from the findings. It proposes a visible safeguards scheme anchored in pre-entry summaries of rules, user rights and redress pathways, supported by a reference safety-UX baseline that platforms can adopt without undermining openness. For accountability, it recommends privacy-preserving credentials (i.e. verifiable credentials and decentralised identifiers with selective disclosure) so that users can prove what's necessary without revealing who they are. To make safety auditable rather than aspirational, it introduces machine-readable safety manifests and explainable moderation trails that regulators, procurers and researchers can verify. Each lever is framed with minimal viable specifications, pragmatic integration routes for existing platforms, and measurement hooks to verify impact in pilots. Finally, the chapter points to procurement clauses and cross-border sandboxing as the engines for scale, turning pilot evidence into market pull while avoiding fragmentation across jurisdictions.

Chapter 6: Conclusions. The report integrates the methodological and substantive strands. It reiterates how a live, factorial design embedded in real platforms produced decision-useful evidence, restates the core signal that visible safeguards increase trust and comfort without harming openness, and clarifies why accountability should be engineered via privacy-preserving credentials rather than identity exposure. It closes by charting the path from evidence to roadmap and from there to standards, procurement and practice in European virtual worlds.

Annexes. **Annex I** details the selected worlds and implementations (layouts, assets, governance/identity cues, fidelity parameters, interaction scripts, configuration notes); **Annex II** provides experiments' material; **Annex III** includes questionnaires and evaluation scales with item wording, response formats, scoring keys and reliability statistics; **Annex IV** reproduces the Ethics Advisor's assessment (clearance, mitigation, recommendations); **Annex V** summarises the DPIA (processing description, risks, flows, retention, governance).

1 BASELINE AND OBJECTIVES

This chapter sets the evidentiary baseline and the concrete objectives that guide our experiment-to-policy pipeline. It explains why the study focuses on governance signalling and digital identity as the pivotal contextual levers in virtual worlds, how these levers are operationalised in two mainstream platforms under tightly controlled technical conditions, and which outcome constructs are used to translate foresight questions into testable hypotheses and interpretable effect sizes. Section 3.1 details the key variables (i.e. Rule Framework and Digital Identity) derived from the OPENVERSE foresight track and aligned with the EU vision of “open but safe” virtual worlds; it also introduces two orthogonal technical stressors, visual fidelity and interaction complexity, to probe ecological sensitivity while keeping hardware, session structure, and locomotion constant. Section 3.2 defines the benchmarking perimeter across technological, ethical, legal, and equality domains so that findings are portable to standards, procurement, and roadmapping, with explicit attention to comfort and accessibility as first-class equity concerns. Section 3.3 documents the two selected worlds and the design choices that make governance and identity cues salient without altering core mechanics, thereby preserving internal validity while maximising external face validity. Together, these elements establish a coherent baseline: a 2x2 factorial design embedded in real platforms, policy-relevant outcomes (trust/privacy, openness, presence, sickness), and a benchmarking frame that allows results to feed directly into T2.5 roadmapping, WP6 standardisation/exploitation, and WP5 cross-initiative alignment.

1.1 KEY VARIABLES

The study centred on two contextual levers derived from the OPENVERSE foresight axes and from recent EU policy frames for virtual worlds: a Rule Framework factor and a Digital Identity Policy factor. The Rule Framework contrasted a platform-only moderation setting where cues pointed to community standards and Terms of Service enforced by the provider with an externally signalled, institution-anchored setting where signage and assets referenced oversight beyond the platform itself. This manipulation operationalises the policy question of whether visible, institutional safeguards can increase trust and comfort without eroding perceived openness, an idea consonant with the European Commission’s emphasis on “open but safe” Web 4.0 trajectories (European Commission, 2023) and with OPENVERSE foresight on “sovereign openness.”

The **Digital Identity Policy** was orthogonal to rules and varied **pseudonymous/anonymous** participation against a **verified/real-name cue** linked to an institutional authenticator. This tests whether accountability achieved through identity salience actually raises trust, or instead dampens immersion, relative to privacy-preserving pseudonymity, an issue flagged both in governance debates and in platform practice (VRChat Documentation, 2022; Veltri, Mureddu, Innocenti, & Sirizzotti, 2025).

Two **technical variables** were included as exploratory, orthogonal stressors to probe ecological sensitivity without confounding locomotion or device pipelines: **visual fidelity** of the world (low vs. high) and **interaction complexity** (simple one-step actions vs. short multi-step sequences). These factors allowed us to check whether UX “heaviness” amplifies or dampens the governance and identity effects while keeping hardware, headset, and session structure constant (see ANNEX II for parametrisation details).

Outcomes were organised in two tiers aligned with policy-relevant endpoints. **Primary outcomes** captured **perceived trust/privacy** and **perceived openness/sustainable innovation** constructs that link directly to EU strategy on trustworthy, human-centric virtual worlds and to OPENVERSE trends work (European Commission, 2023; OPENVERSE D3.1, 2024).

Secondary outcomes addressed experiential quality and safety (see Annex III) using a short **Igroup Presence Questionnaire (IPQ)** subset for spatial presence/immersion and the **Simulator Sickness Questionnaire (SSQ)** for comfort, with optional emotional/social presence items and **behavioural indicators** from system logs when available; these measures also reflect the broader safety- and privacy-by-design discourse in international policy fora (World Economic Forum, 2023; XR Safety Initiative, 2020–2022).

To reduce compositional confounds, **VR literacy** prior exposure to immersive tech and gaming was measured at baseline and used as a **covariate**, so observed differences could be attributed to the manipulations rather than familiarity effects.

Taken together, the variables instantiate a **2×2 factorial design** (Rule Framework × Digital Identity) embedded in two real platforms and enriched by orthogonal UX stressors. This design translates foresight questions into testable contrasts with interpretable effect sizes on trust/privacy, openness, presence, and sickness, thus providing an evidence bridge from scenario logic to roadmapping and standardisation (European Commission, 2023; OPENVERSE D3.1, 2024).

1.2 AREA OF BENCHMARKING (TECHNOLOGICAL, ETHICAL, LEGAL, EQUALITY)

The benchmark spans four interlocking domains to ensure that evidence from the exploratory study can travel into standards, procurement, and policy design. On the technological side, we benchmark platform affordances and user-experience determinants that are known to shape trust and comfort: headset/runtime stability on Meta Quest; scene customisability for governance and identity cues; and two orthogonal UX stressors such as visual fidelity and interaction complexity, varied where feasible to test ecological sensitivity without confounding locomotion or rendering pipelines. Core outcome instrumentation couples experiential quality (presence/immersion via a short IPQ subset) with safety and accessibility (Simulator Sickness Questionnaire, SSQ), so that differences observed between Horizon Worlds (externally signalled rules) and VRChat (platform-only norms) can be attributed to governance and identity treatments rather than to hardware idiosyncrasies. This technical frame enables clean, comparable effect estimates across the 2×2 factors (Rule Framework × Digital Identity) and across exploratory UX conditions (fidelity, complexity).

Ethical benchmarking focuses on dignity, comfort, and informed participation. The protocol implements ex-ante consent with explicit right to withdraw; a brief acclimatisation to reduce novelty stress; immediate post-exposure debrief; and monitoring for discomfort, with SSQ as a quantitative guardrail. Instrumentation choices and data-cleaning practices (dropping invariant items; recoding reverse-keyed items) are documented to avoid artefactual reliability gains and to preserve construct validity, with presence and baseline-privacy flagged for refinement in subsequent waves. We also treat identity exposure as an ethical variable in its own right measuring whether verified identity cues alter immersion or perceived safety, so the benchmark can distinguish benefits of accountability from potential harms to autonomy and self-expression.

The legal dimension benchmarks alignment with the EU's prevailing digital-policy stack and its implications for virtual worlds. While this pilot does not adjudicate compliance, it operationalises constructs that are policy-salient - trust/privacy, openness, safety-by-design - and stages them in environments that model externally signalled governance versus platform self-regulation, echoing directions in European foresight and roadmapping (WP2/T2.4 and T2.5). Abbreviations and scoping in the deliverable template reflect the regulatory landscape (e.g., GDPR for data protection; DSA/DMA for platform duties; eIDAS for identity assurances), and the experiment's identity arm explicitly informs the trade-offs between real-name exposure and privacy-preserving attestations to be taken up by T2.5. In this sense, legal benchmarking is evidence-led: we test whether levers often assumed to raise trust

(rule visibility; identity salience) in fact do so, and we quantify any UX costs that policymaking should internalise.

Finally, equality benchmarking addresses inclusion, accessibility, and non-discrimination across the study flow and measurement. The sample is gender-balanced at the cell level (overall 56% female across conditions), randomisation is balanced to avoid compositional biases, and VR literacy is captured as a covariate so that differences in prior exposure do not masquerade as treatment effects. Comfort and sickness are treated as first-class equity variables because they disproportionately affect who can participate in immersive experiences; reductions in sickness under externally signalled rules therefore count as accessibility gains as well as safety gains. The use of pseudonymous identity conditions also probes whether anonymity supports participation and presence for users who might otherwise be marginalised by real-name norms. Together, these equality checks make the benchmark sensitive to who benefits or is excluded when governance and identity levers are pulled.

1.3 SELECTED VIRTUAL WORLD

For the exploratory study we purposefully selected two widely used social VR platforms, namely Meta's Horizon Worlds and VRChat (see ANNEX I) because, taken together, they span the governance and identity archetypes required by the experimental design and can be deployed reliably on Meta Quest headsets. Horizon Worlds offers a tightly signalled, institution-anchored governance style that served as a proxy for an externally framed rule context; VRChat, by contrast, is a creator-driven ecosystem with strong user-generated content and long-standing pseudonymous participation, which we used as a proxy for a platform-moderated or "wild-west" rule context. This pairing was not about branding the platforms as "good" or "bad," but about operationalising two ends of the foresight axis such as visible institutional safeguards versus self-regulatory norms under comparable technical conditions so that user-level outcomes could be measured cleanly.

Horizon Worlds was configured to make **external rule visibility** salient at the point of entry and within the scene. Participants encountered environmental cues (e.g., signage and visual assets) that referenced an institutional presence and formal oversight, consistent with a world where rules are set or audited by authorities beyond the platform itself. This environment therefore embodied the "council-regulated" logic used in our factorial design and allowed us to test whether explicit governance signalling raises perceived trust and comfort without depressing openness. The same Horizon scene also accommodated the **identity-verified** condition by displaying a visible verification marker linked to an institutional authenticator, thereby permitting us to examine immersion and comfort when identity salience is increased. Implementation choices were constrained to cosmetic and informational cues—no change in technical locomotion or device settings, so that any measured differences would track governance signalling rather than hardware or rendering changes.

VRChat was selected to represent a **user-driven, rules-light ambience** in which platform terms and community norms remain present but are less externally signalled. The chosen world emphasised creator assets and standard community ruleboards rather than institutional branding, making **platform-only governance** the dominant cue. VRChat's longstanding support for pseudonymous play and its Trust/Safety affordances gave us a realistic baseline for the **anonymous identity** condition, while still allowing a visible "verified" cue to be introduced for the contrasting identity treatment. As with Horizon, the VRChat configuration preserved parity in movement, interaction basics, and session length; the design intent was that any divergence in outcomes would reflect differences in perceived governance and identity posture rather than differences in world mechanics.

Together, the two platforms enabled the **2x2 structure** of the study: Rule Framework (external/institutional vs. platform-only) crossed with Digital Identity Policy (verified/real-name cue vs. pseudonymous/anonymous cue). The same pair also accommodated two **exploratory technical**

factors, visual fidelity (lower-fidelity, stylised assets versus higher-fidelity textures) and **interaction complexity** (simple one-step interactions versus short multi-step sequences), so that we could probe whether UX “heaviness” amplifies or dampens the governance and identity effects. These technical factors were held orthogonal to the main manipulations and varied only when feasible, with headset, session timing, and onboarding kept constant across all conditions.

The selection also met practical criteria important for a pilot: broad **device accessibility** (Quest-native support), stable **session orchestration** for individual participants, and sufficient **scene customisability** to embed rule and identity cues without altering core locomotion or rendering pipelines. By anchoring the experiment in two mature ecosystems that are widely recognised by users and policymakers, the study enhances the external face validity of its comparisons while keeping internal validity through tightly controlled cue design and exposure. Full implementation details, including the precise assets used for governance signalling and identity markers, the scene layout, and the parameterisation of visual fidelity and interaction steps are provided in **ANNEX I** for both selected worlds.

2 METHODOLOGY OF THE EXPERIMENT

This chapter describes the experimental protocol used to generate user-level evidence on how governance cues and digital identity policies shape perceptions of trust, privacy, and openness in virtual worlds, with a view to informing roadmapping activities of T2.5. Building on the scenario work of D2.3 Technological Foresight (T2.4) and the project's policy aims around openness, safety, and human-centric innovation, we implemented a 2x2 factorial design contrasting a Rule Framework factor, platform-only moderation versus externally signalled, institution-anchored rules, and a Digital Identity Policy factor, pseudonymous/anonymous use versus verified/real-name cues. Two technical factors were included as exploratory, orthogonal dimensions to probe ecological sensitivity: visual fidelity (low vs high) and interaction complexity (simple vs complex). The study recruited a convenience sample of university students, ultimately N=101, who completed baseline questionnaires, a brief acclimatisation, a short-guided exploration in the assigned condition using Meta Quest headsets in Horizon Worlds and VRChat, and a post-experience battery capturing the primary outcomes and secondary experiential and safety measures. The measures and analysis plan were chosen to translate scenario and policy questions into testable hypotheses and interpretable effect sizes, enabling both within-subject estimates of learning/onboarding and between-group tests of the governance and identity manipulations.

2.1 PARTICIPANTS

The study enrolled a convenience sample of university students with a planned target of about fifty; in practice, one-hundred-and-one participants completed the full pre/post protocol and were randomized evenly across the four cells of the 2x2 design (council-regulated with identity: 25; council-regulated without identity: 25; wild-west with identity: 26; wild-west without identity: 25).

This cohort was selected for pragmatic and methodological reasons: ready accessibility within an ethics-approved setting, broadly comparable levels of digital/VR literacy to support clean construct testing, and efficient recruitment for a tightly controlled pilot intended to isolate governance and identity effects before scaling to more heterogeneous populations.

The sample skewed slightly female overall (56%), with comparable gender shares in each cell, supporting balance across conditions. Participants were between 18 and 28 years old, gave informed consent, and completed baseline instruments before any headset exposure. These included demographics and a "VR literacy" questionnaire capturing prior use of immersive devices, gaming, and virtual environments; VR literacy was later used as a covariate to guard against familiarity effects when estimating treatment impacts on outcomes. A short familiarization phase preceded the main exposure, after which participants filled post-experience batteries; data screening removed a small set of invariant items and recoded reverse-keyed items, leaving no material missingness for analysis. This combination of age-bounded student cohort, balanced randomization, documented consent and acclimatization, explicit capture of prior experience, and a clean post-processing pipeline and provides a transparent, policy-relevant baseline for interpreting effects on trust, privacy, openness, presence, and comfort.

2.2 MATERIALS AND APPARATUS

Sessions were conducted with Meta Quest VR headsets, and the virtual environments were created within Meta's Horizon Worlds and VRChat. These two platforms were chosen for their contrasting environments: Horizon Worlds, a Meta-operated platform with ties to Facebook/Meta account policies (proxy for a more institutionally-controlled environment), and VRChat, an independent platform known for user-generated content and relative anonymity (proxy for a more user-driven or "wild-west"

environment). Within these platforms, custom virtual scenarios were set up to manipulate two primary factors as summarized in the following figure:




Figure 1: The two dimensions of the experiments and the 4 quadrants.

- Rule Framework:** Participants were randomly assigned to either a platform-only condition or an external-rule condition. In the platform-only condition (implemented in VRChat), the virtual world contained cues emphasizing only the platform's own rules/Terms of Service and community moderation norms. This reflects a scenario where control is left to the platform provider (minimal outside intervention). In the external-rule condition (implemented in Horizon Worlds), the environment included cues highlighting external legal frameworks and an institutional presence, for example, signs referencing a European trust authority or public agency monitoring conduct. This simulates a scenario where an external public authority imposes rules in the virtual world space.
- Digital Identity Policy:** Orthogonal to the above, the second factor was identity disclosure vs. anonymity. In the verified identity condition, participants' avatars were tagged or presented as being verified by an institution (e.g. displaying a badge of authentication, indicating the user's identity had been verified through a government or university system). In practical terms, participants in this condition might have been assigned an avatar or profile name that signaled a "real" identity (for instance, a full name or an institutional ID). In the anonymous identity condition, participants used a generic avatar/name with no personally identifying information, akin to a default or pseudonymous account. This factor tests the effect of real-name/verified identity requirements (as advocated in some safety-oriented scenarios) against pseudonymity (as in open or privacy-preserving scenarios).

Additionally, the study noted **two technical factors** such as visual fidelity of the environment and interaction complexity (simple vs. complex tasks) which were varied in some trials as exploratory variables (see the following figure). These were not core to the hypotheses but included to observe any influence on user experience (for example, whether a more cartoonish world vs. a realistic world alters trust or presence). All participants experienced the simulation individually.




Figure 2: 3D space for Two technical factors and the Two main dimensions of the analysis

2.3 PROCEDURE

At the start of each session, participants completed baseline questionnaires on demographics, VR literacy, and attitudes toward privacy/trust. They were then randomly assigned to one of the four experimental groups in a 2 (Rule Framework: platform vs. external) $\times$ 2 (Identity: anonymous vs. verified) factorial design. Importantly, assignment was balanced so that roughly equal numbers experienced each condition (approximately 25 per cell, ultimately, e.g. ~25 in “external rules + identity”, ~25 in “external rules + no identity”, etc.). Participants donned the VR headset and underwent a short familiarization phase (~15 minutes) in the virtual environment to practice navigation and interaction basics. This acclimation was to minimize novelty effects and motion sickness.

Next came the exploration phase (~10–15 minutes) wherein the participant engaged in the virtual world corresponding to their assigned condition. During this time, they could interact with virtual objects or scripted scenarios and encounter the cues relevant to the experimental manipulations. For example, a participant in the external-rule/verified ID condition (Horizon Worlds scene) might see a virtual “EU Trust Authority” booth and their avatar might have a visible real-name label; whereas one in platform-only/anonymous (VRChat scene) might see user-generated posters of VRChat community rules and interact with others under a pseudonym. The interactions were kept light to ensure all participants had comparable exposure e.g. exploring a virtual public plaza, reading displayed messages, possibly greeting a confederate avatar. Researchers monitored sessions to ensure consistency (and to observe any notable user reactions or issues).

After the exploration, participants removed the headset and immediately filled out a post-experience questionnaire battery. This survey gathered the primary outcome measures: perceived trust and privacy in the environment, and perceived openness and innovation in the platform. These constructs were measured via multi-item Likert scales developed for the project (including a combined trust/privacy scale and an openness scale). The emphasis on “openness and sustainable innovation” aligns with foresight dimensions identified in Task 2.4 essentially gauging if the environment felt open, transparent and supportive of innovation (key values in the European vision). Secondary outcomes included standard questionnaires on presence/immersion (e.g. Igroup Presence Questionnaire, IPQ), user comfort (including a Simulator Sickness Questionnaire to check for motion sickness symptoms), and any emotional or social presence scales if time allowed. Participants could also provide qualitative feedback about their experience, especially regarding their sense of safety, identity, and any concerns. Finally, a debriefing was conducted where the true purpose of the study was explained (that the focus was on differences in platform rules and identity policies) and participants could discuss any

observations or discomfort. Ethically, they were reminded of their right to withdraw their data if they had any concerns (none did).

2.4 MEASURES

Primary outcomes measured with the experiment were perceived trust/privacy and perceived openness, two constructs chosen to operationalise the OPENVERSE foresight axes on openness and sustainable, human-centric innovation (see D2.3 of OPENVERSE). These outcomes align with EU policy priorities for Web 4.0 and virtual worlds, notably the Commission's strategy on *openness, safety, and fairness* and the emphasis on trustworthy identity and governance (European Commission, 2023; European Commission, Digital Strategy, 2023/2025). They also speak to the broader international policy conversation on privacy-by-design and safety-by-design for immersive media (World Economic Forum, 2023; XR Safety Initiative, 2020–2022) and to academic/legal analyses of identity, privacy and safeguards in virtual worlds (Wu, 2023; Polychronaki et al., 2024; Laiz et al., 2024). Secondary outcomes captured experiential quality and safety, spatial presence/immersion (short IPQ battery), user comfort via the Simulator Sickness Questionnaire (SSQ), plus optional emotional and social presence short forms when time permitted; exploratory behavioural indicators (e.g., interaction counts) were derived from platform/system logs (see also OPENVERSE D3.1 on *trends and governance factors*). Together, these measures were intended to create a direct bridge from user-level evidence to policy-relevant endpoints for roadmapping and standards (European Commission, 2023; OPENVERSE D3.1, 2024; European Commission, Digital Strategy, 2023/2025), so to inform T2.5 policy recommendations as well as WP5 (Collaboration and Synergies) and WP6 Impact creation.

The analytic dataset comprised $N = 101$ participants with pre- and post-test surveys and no material missingness after cleaning; randomisation yielded balanced 2×2 cells across Rule Framework and Digital Identity conditions. Sessions were run on Meta Quest headsets within Horizon Worlds (proxy for externally framed/institutionally signalled rule contexts) and VRChat (proxy for platform-moderated, user-driven contexts with established safety/trust affordances such as trust ranks), consistent with the study's aim to contrast governance/identity approaches found in real platforms (Veltri, Mureddu, Innocenti, & Sirizzotti, 2025; VRChat Documentation, 2022). The procedure included a brief familiarisation phase ($\leq 15'$), a 10–15' exploration under condition-specific cues, and a post-experiment battery covering trust/privacy, openness, presence/immersion, comfort/SSQ, and optional emotional/social presence items.

Data quality controls followed standard psychometric practice prior to scale construction. First, zero-variance items (nine IPQ items; three baseline-privacy items) were removed because they cannot discriminate and can artificially inflate reliability; second, all reverse-keyed items (IPQ items 5/6/10/12; baseline-privacy 1/4/6) were recoded so higher scores uniformly indicated more of the construct. Reliability (Cronbach's α) on retained items showed good internal consistency for Openness ($\alpha \approx .77$) and excellent consistency for SSQ ($\alpha \approx .91$); Trust/Privacy was acceptable ($\alpha \approx .63$), and VR literacy (pre-test covariate) was good ($\alpha \approx .84$). Two scales were flagged as developmental in this pilot: Presence was low ($\alpha \approx .24$) after informative-item retention, and Baseline privacy attitudes became unreliable ($\alpha \approx -.09$) once invariant items were dropped. As a result, cleaned composites were retained for exploratory analyses (primary endpoints: trust/privacy, openness; secondary endpoints: presence/immersion, SSQ), while Presence and Baseline privacy are earmarked for instrument refinement in subsequent waves (e.g., rewording, discriminative items, or alternative short forms).

2.5 STATISTICAL MEASURES

Analyses combined within-subject change estimation with between-group inference to evaluate how Rule Framework (external/institutional vs. platform-only) and Digital Identity (verified vs. anonymous) shape the primary outcomes (trust/privacy, openness) and secondary outcomes (presence/immersion,

comfort via SSQ), with VR literacy entered as a covariate. Following the analysis plan, we first quantified the education/onboarding effect using paired-samples t-tests on pre → post trust/privacy for the full sample (two-tailed $\alpha=.05$), reporting Cohen's d for dependent means with 95% CIs (Hedges' g where small-sample bias correction is desired). We then fitted two-way ANOVA/ANCOVA models for each post-test endpoint with fixed factors Rule Framework and Digital Identity, VR literacy as a covariate, and the Rule $\times$ Identity interaction, reporting omnibus F tests, partial η^2 effect sizes, and model-based means with CI bands. These models implement the core comparisons specified in the protocol and used for the main results (e.g., trust/privacy main effect of Rule Framework; presence main effect of Identity; SSQ reduction under external rules).

To increase precision and leverage the pre/post structure where applicable, we estimated complementary mixed-effects models (restricted maximum likelihood) with Participant as a random intercept and Time (pre, post) as a within-subject fixed factor crossed with Rule Framework and Digital Identity as between-subject factors. This structure provides a check on the paired-t and ANCOVA conclusions and accommodates any residual heterogeneity across participants (model form for trust/privacy: $Score \sim Time \times Rule \times Identity + VR\text{-}literacy + (1|Participant)$). Where the exploratory technical factors (Visual Fidelity; Interaction Complexity) were manipulated, they were added as fixed effects (and, if cell sizes allowed, interacted with Rule or Identity) but treated as exploratory with appropriately cautious interpretation:

- Assumptions and robustness.** Model diagnostics included Shapiro-Wilk tests and Q–Q plots of residuals (normality), Levene's tests (homogeneity), and checks for influential cases ($| \text{standardised residual} | \geq 3$). Where ANOVA assumptions were materially violated, we report Welch's corrections (unequal variances) and/or HC3 heteroskedasticity-robust standard errors. For mixed models, residual distributions and random-effect variance components were inspected; if necessary, we applied rank-based or log transformations to stabilise variance for SSQ. Given the limited number of primary endpoints, two-sided $\alpha=.05$ was retained; for families of secondary/exploratory tests we applied Benjamini–Hochberg false-discovery-rate control to limit Type-I error inflation. All tests report effect sizes (partial η^2 , generalised η^2 where repeated-measures are prominent; Cohen's d for mean changes) with 95% CIs to support interpretation beyond p -values.
- Correlational structure and derived metrics.** To characterise how experiential variables covary, we computed Pearson correlations among trust/privacy, openness, presence, and SSQ (Spearman's ρ as a robustness check for non-normal pairs). These matrices illuminate the observed pattern whereby higher presence mildly co-occurs with more sickness and slightly lower trust, whereas trust/privacy correlates positively with openness, a relationship of policy relevance for openness-with-safeguards strategies in EU Web 4.0 initiatives (European Commission, 2023; European Commission, Digital Strategy, 2023/2025). For transparency, we report sample size per correlation, two-tailed p , and 95% CIs.
- Missing data and outliers.** As noted in the data-quality section, no material missingness remained after cleaning; analyses used listwise inclusion. Outliers were handled conservatively: we flagged observations exceeding ± 3 SD on standardised residuals and verified that substantive conclusions were unchanged by their exclusion (leave-one-out sensitivity). Scale construction steps (removal of zero-variance items; reverse-key recoding) preceded all analyses to avoid reliability artefacts, with internal consistency reported for each composite (e.g., Openness $\alpha \approx .77$, Trust/Privacy $\alpha \approx .63$, SSQ $\alpha \approx .91$), and Presence/Baseline privacy treated as developmental measures (low/negative α) and interpreted cautiously.
- Interpretive framing.** Finally, we read experimental effects against the consortium's policy context—i.e., whether externally signalled rule frameworks can raise trust and comfort

without depressing perceived openness, and whether identity verification (vs. pseudonymity) adds trust or impairs immersion—to align statistical conclusions with foresight and roadmapping aims in OPENVERSE and with ongoing EU strategy on virtual worlds (European Commission, 2023; European Commission, Digital Strategy, 2023/2025; OPENVERSE D3.1). Platform-specific affordances (e.g., VRChat trust/safety ranks) inform interpretation of identity cues as presently implemented (VRChat, 2022), while broader guidance on privacy- and safety-by-design in immersive media contextualises our secondary outcomes (WEF, 2023).

3 RESULTS OF THE EXPERIMENTS

As already discussed in the previous Chapter, a total of 101 participants completed both pre- and post-test measures in this virtual reality experiment examining the effects of regulatory framework and identity policy on user experience. The sample was well-balanced across experimental conditions, with 25 participants in the council-regulated with identity disclosure condition (Vrc-id), 25 in the council-regulated without identity disclosure condition (Vrc-noid), 26 in the wild-west with identity disclosure condition (W-id), and 25 in the wild-west without identity disclosure condition (W-noid). The overall sample composition was 56% female, with consistent gender distribution across all experimental conditions (ranging from 56% to 58% female across groups).

3.1 SCALE RELIABILITY AND DATA CLEANING

Prior to conducting the main analyses, comprehensive data screening was performed to ensure the integrity and reliability of all measurement instruments. This process involved systematic examination of item-level characteristics and correction of methodological artefacts that could compromise the validity of subsequent analyses.

Two critical issues were identified and addressed during the data cleaning process. First, twelve questionnaire items demonstrated zero variance, with all 101 participants providing identical responses. This included nine items from the Igroup Presence Questionnaire (IPQ) and three items from the baseline privacy scale. Items with zero variance cannot discriminate among participants, fail to correlate with other variables, and artificially inflate reliability estimates while adding only a constant to total scores. These items were therefore excluded from both reliability calculations and scale score computations.

Second, reverse-keyed items required recoding to ensure consistent directionality across all measures. Four IPQ items (items 5, 6, 10, and 12) and three baseline privacy items (items 1, 4, and 6) were negatively worded and required response scale inversion (e.g., transforming responses of 1 → 7, 2 → 6) so that higher values consistently indicated greater levels of the target construct. Following these data cleaning procedures, reliability analyses were conducted for all scales using Cronbach's alpha. The results revealed substantial variation in internal consistency across measures, as presented in Table 1.

Table 1: Scale Reliability and Item Retention

Scale	Items Retained	Cronbach's α	Reliability Assessment
VR Literacy (pre-test)	6/6	.84	Good
Baseline Privacy (pre-test)	3/6	-.09	Unreliable
Trust/Privacy (post-test)	6/6	.63	Acceptable
Openness (post-test)	6/6	.77	Good
Presence (post-test)	4/13	.24	Low
Simulator Sickness (post-test)	16/16	.91	Excellent

The VR literacy scale demonstrated good reliability ($\alpha = .84$) with all six original items retained. The simulator sickness questionnaire showed excellent internal consistency ($\alpha = .91$) with all 16 items contributing meaningfully to the scale. The openness scale demonstrated good reliability ($\alpha = .77$), while the trust/privacy scale achieved acceptable levels of reliability ($\alpha = .63$).

However, two scales showed concerning reliability patterns. The baseline privacy scale, after removing zero-variance items, yielded a negative alpha coefficient ($\alpha = -.09$), indicating that the remaining three items did not form a coherent construct in this sample. Similarly, the presence scale reliability dropped substantially from the original 13 items to just four informative items, resulting in low internal consistency ($\alpha = .24$). These low reliability values suggest that the pilot sample demonstrated limited variability on these particular constructs. However, the cleaned scores remain statistically defensible and free from constant-item artefacts.

3.2 PRE-POST CHANGES IN TRUST AND PRIVACY PERCEPTIONS AND EFFECTS OF MANIPULATIONS

To examine whether the virtual reality experience influenced participants' trust and privacy perceptions, a paired-samples t-test was conducted comparing pre-test and post-test scores on the trust/privacy measure. The analysis revealed a statistically significant increase in perceived trust and privacy following the VR experience, $t(100) = -3.70$, $p < .001$, Cohen's $d = 0.53$. This represents a medium-sized effect, indicating that participants generally reported enhanced trust and privacy perceptions after engaging with the virtual environment, regardless of experimental condition.

The primary research questions focused on the differential effects of regulatory frameworks (council-regulated vs. wild-west) and identity policies (identity disclosure vs. anonymous interaction) on various aspects of the VR experience, as introduced in D2.3 on technological foresight. These effects were examined using a series of 2×2 analysis of covariance (ANCOVA) models, with VR literacy included as a covariate to control for individual differences in virtual reality familiarity and expertise.

The ANCOVA examining trust and privacy perceptions revealed a significant main effect of regulatory framework, $F(1, 96) = 5.02$, $p = .027$. Participants in council-regulated virtual worlds reported significantly higher levels of trust and privacy compared to those in wild-west environments. The identity policy manipulation did not produce a significant main effect, $F(1, 96) = 0.25$, $p = .618$, nor was there a significant interaction between regulatory framework and identity policy, $F(1, 96) = 0.10$, $p = .754$. The VR literacy covariate was not significantly related to trust and privacy outcomes, $F(1, 96) = 0.05$, $p = .828$.

Analysis of openness and sustainability attitudes showed no significant effects of either experimental manipulation. The regulatory framework main effect was not significant, $F(1, 96) = 1.87$, $p = .175$, nor was the identity policy main effect, $F(1, 96) = 0.71$, $p = .401$. The interaction between these factors was also non-significant, $F(1, 96) = 0.01$, $p = .938$. VR literacy did not significantly predict openness attitudes, $F(1, 96) = 1.37$, $p = .245$.

The analysis of simulator sickness revealed a highly significant main effect of regulatory framework, $F(1, 96) = 10.70$, $p = .001$. Participants in council-regulated environments experienced substantially lower levels of simulator sickness compared to those in wild-west conditions. This effect was particularly notable given its large magnitude, with council-regulated worlds reducing simulator sickness scores by approximately half relative to unregulated environments. Neither the identity policy main effect, $F(1, 96) = 1.29$, $p = .259$, nor the regulatory framework $\times$ identity policy interaction, $F(1, 96) = 0.34$, $p = .560$, reached statistical significance. VR literacy was unrelated to simulator sickness, $F(1, 96) = 0.02$, $p = .884$.

The spatial presence analysis yielded a significant main effect of identity policy, $F(1, 96) = 4.85$, $p = .030$. Participants required to disclose their real identity reported lower levels of spatial presence compared to those who could interact anonymously. This effect represented approximately one-quarter of a standard deviation decrease in presence when real identity disclosure was required. The main effect of the regulatory framework was not significant, $F(1, 96) = 0.61$, $p = .438$, nor was the interaction between the regulatory framework and identity policy, $F(1, 96) = 0.69$, $p = .408$. VR literacy did not significantly predict spatial presence, $F(1, 96) = 0.41$, $p = .521$.

3.3 CORRELATIONAL PATTERNS AMONG OUTCOME VARIABLES

Zero-order correlations among the primary outcome variables revealed several meaningful associations that provide insight into the relationships between different aspects of the VR experience (Table 2).

Table 2: Intercorrelations Among Primary Outcome Variables

Variable	1	2	3	4	5
1. VR Literacy	-	.02	.11	.07	.00
2. Trust/Privacy		-	.61***	-.26**	-.43***
3. Openness			-	-.09	-.14
4. Presence				-	.32**
5. Simulator Sickness					-

*Note: $p < .01$, * $p < .001$ *

The strongest correlation emerged between trust/privacy and openness ($r = .61$, $p < .001$), suggesting that participants who felt more secure and trusted in the virtual environment were also more open to sustainability-related concepts and behaviours. Trust and privacy perceptions were negatively correlated with simulator sickness ($r = -.43$, $p < .001$), indicating that participants who experienced greater comfort and security in the virtual environment were less likely to experience adverse physiological symptoms.

Interestingly, spatial presence showed a positive correlation with simulator sickness ($r = .32$, $p < .01$), suggesting that more immersive experiences may come at the cost of increased physiological discomfort. Presence was negatively correlated with trust and privacy perceptions ($r = -.26$, $p < .01$), indicating a potential trade-off between immersive engagement and feelings of security within the virtual environment.

VR literacy showed minimal correlations with all outcome variables, with correlation coefficients ranging from .00 to .11, none of which reached statistical significance. This pattern is consistent with the ANCOVA results showing that VR literacy was not a significant predictor in any of the experimental models.

The experimental manipulation of the regulatory framework emerged as the most consistent and impactful factor influencing the quality of the VR experience. Council-regulated virtual worlds produced two major benefits: significantly enhanced trust and privacy perceptions and dramatically reduced simulator sickness compared to wild-west environments. The magnitude of the simulator sickness reduction was particularly striking, with regulated environments cutting adverse symptoms by approximately half.

The identity policy manipulation showed a more selective pattern of effects. While requiring real identity disclosure did not influence trust, openness, or simulator sickness, it did produce a reliable decrease in spatial presence. This suggests that anonymity may facilitate deeper immersive engagement, even when it does not affect other aspects of user experience or comfort.

Notably, no significant interactions emerged between regulatory framework and identity policy across any outcome measure, indicating that these factors operate independently rather than synergistically. The VR literacy covariate consistently failed to predict any outcome variables, suggesting that individual differences in virtual reality experience and expertise did not meaningfully influence responses to the experimental manipulations in this sample.

The correlational analyses revealed a coherent pattern of relationships among outcome variables, with trust and privacy serving as a central hub connecting to multiple other aspects of the VR experience. The strong positive correlation between trust and openness suggests that creating secure virtual environments may facilitate greater receptivity to educational or persuasive content. Conversely, the negative correlations between trust and both simulator sickness and spatial presence highlight potential tensions in VR design, where maximising immersion or minimising discomfort may require careful balance with security and privacy considerations.

4 BENCHMARK CONCLUSIONS

The exploratory experiment demonstrates that modest, well-signalled governance cues can lift **trust/privacy** perceptions and **reduce sickness**, without depressing perceived **openness**; by contrast, **verified identity** markers did **not** add trust or openness and tended to **lower immersion**. Correlational patterns suggest **user safety/comfort** is a hinge variable: when comfort improves, trust tends to rise, and vice versa. These empirical signals map cleanly onto OPENVERSE foresight priorities, especially the need to balance **openness** with **institutional safeguards**, and they point to specific **trust tools** (labels, privacy-preserving credentials, user-agency mechanisms, and sandboxes) that the EU can deploy to scale those benefits.

4.1 HOW THE EXPERIMENTAL RESULTS VALIDATE (AND NUANCE) THE FORESIGHT TRACK IN D2.3

Here we translate the experiment's empirical signals into the language of **D2.3 Technological Foresight**. The aim is twofold: first, to **validate** core assumptions behind the preferential track (that visible, institutional safeguards can raise trust and comfort **without** suppressing openness); second, to **nuance** debates where evidence points to better levers (e.g., privacy-preserving credentials rather than real-name mandates). The three short subsections below organise the findings around (i) **governance signalling**, (ii) **identity policy**, and (iii) **user comfort/UX** the practical hinges through which Europe can pursue "**sovereign openness**" in virtual worlds. Takeaway from the experiments results in relation to the D2.3 can be summarized as follows

- (A). **Institutional safeguards without "closing" the world.** The **external-rule** condition produced **higher trust/privacy** and **much lower simulator sickness** than laissez-faire settings, yet **openness** perceptions stayed statistically unchanged across conditions. This aligns with D2.3's call for "**ambidextrous**" European pathways that combine **openness** with **governance capacity**, i.e., the report's logic of a **preferential track** where institutional signalling coexists with innovation and human-centric design. In D2.3 language, these findings operationalise "**sovereign openness**": external rule visibility can raise safety/trust outcomes without foreclosing creative latitude.
- (B). **Identity visibility vs. pseudonymous accountability.** The **verified identity** cue did **not** raise trust or openness and **reduced presence**; **anonymity** preserved immersion **without** detectable safety losses. This nuances D2.3's governance debates by indicating that **hard real-name policies** are not a necessary condition for perceived safety/trust gains, especially if rule frameworks and safety affordances are well signalled. The foresight implication is to **de-couple safety from real-name identity** and prioritise **assurance mechanisms** that are **portable** and **privacy-preserving** rather than identity-exposing.
- (C). **Comfort as a policy-relevant UX lever.** The negative links between **immersion ↔ trust** and positive links **sickness ↔ immersion** suggest that **comfort/safety baselines** are central to user trust in practice. D2.3's meta-scenario emphasises capacity-building and testing of socio-technical guardrails; the data here bolster that emphasis and justify **UX-first safeguards** as part of any "open-but-safe" trajectory.

4.2 BENCHMARK AGAINST "TRUST TOOLS FOR IMMERSIVE TECHNOLOGIES" (WP3)

The **Trust Tools** draft deliverable under development in WP3 and the outcomes from the related questionnaire, proposes concrete levers Europe can deploy **interoperable identity/credentials, trust**

labels/certification, AI audit trails, user agency controls, and regulatory sandboxes, with a staged policy roadmap (2025–2030+). Each maps to an experimental signal:

1. **Trust labels and transparent rule signalling.** Because external rule cues increased **trust** and reduced **sickness without hurting openness**, an **EU Trust Label** that certifies **safety baselines**, **moderation quality**, and **accessibility** should scale these effects across platforms. Labels make governance **legible** at entry, echoing our treatment effect mechanism (visible safeguards → higher perceived safety/comfort). **Design principle:** certification must be **transparent and verifiable** and embedded in choice architectures so users see it **before** they commit to a world.
2. **Interoperable, privacy-preserving identity & credentials.** Since **real-name verification** depressed **immersion** without adding **trust**, the more promising route is **privacy-preserving credentials** (e.g., **selective disclosure** age/eligibility proofs via **DIDs/VCs**) that **signal accountability** without exposing identity. This aligns with the report’s “Interoperable Identity and Credentials” stack and the eIDAS-linked EU wallet vision. **Policy:** expand **eIDAS→DID/VC** hybrids and mutual recognition, favouring **attribute-based** proofs over identity exposure.
3. **Audit trails and explainability for safety-by-design.** Given the strong role of **comfort/sickness** in trust, platforms should document how **AI moderation** and **safety systems** act in-world (tamper-evident **audit trails**, **explainability** for flags/content rankings), enabling user redress and supervisory audits, key ingredients in the Trust Tools blueprint and consistent with the AI Act orientation.
4. **User-agency controls.** To protect immersion and safety, bake in **privacy dashboards**, **dynamic consent managers**, **panic/exit** affordances, **harassment filters**, and **well-being timers**. Our correlations imply such tools could lift trust indirectly by reducing discomfort and perceived risk, matching the brief’s “user empowerment” design principle.
5. **Regulatory sandboxes.** The experiment itself is a micro-sandbox validating that **visible governance** increases **trust** without harming **openness**. Next step: **EU-level sandboxes** to test **labels**, **wallet integrations**, **audit trails**, and **inter-world portability** under a “28th regime,” then fast-track successful practices into certification and procurement.
6. **Roadmap fit.** The Trust Tools report’s staged actions (2025–2027 pilots and label launch; 2027–2030 scaling; 2030+ globalisation) provide a policy spine for scaling what our experiment observed at small-N. The **label + wallet/VC + audit + agency + sandbox** stack is the most direct way to reproduce the **external-rules** gains at ecosystem scale **without** resorting to blunt **real-name** mandates.

4.3 IMPLICATIONS FOR OPENVERSE ROADMAPING (T2.5)

This section turns the experiment’s evidence into **actionable levers** for the T2.5 roadmap. The focus is pragmatic: pick the **highest-yield interventions** that measurably lift trust and comfort **without** eroding openness, and stage them through pilots, certification, and standards. What follows prioritises (i) **governance signalling** users can see, (ii) **privacy-preserving identity** over exposure, (iii) **safety-by-design UX baselines**, (iv) **auditability** for supervision and redress, and (v) **cross-border sandboxes** to prove and scale what works. Each item is framed so it can become a **policy option with KPIs, owners, and timing** in the Roadmap:

- (i). **Prefer “visible safeguards” over “identity exposure.”** The data support **governance signalling** (labels, external rules, safety affordances) as the high-yield lever. Prioritise **rule legibility** and

comfort engineering before debating identity. This concretises D2.3's **sovereign openness**: openness endures when safety is made **explicit, auditable, and portable**.

- (ii). **Adopt privacy-preserving credentials.** For risk-sensitive contexts (age-gating, access control), **attribute proofs** via **VCs/DIDs** should replace blanket real-name exposure. This both preserves **immersion** and sustains **trust**, matching our identity findings and the Trust Tools identity section.
- (iii). **Standardise safety UX baselines.** Given the trust–comfort links, develop a **reference safety UX** (panic/exit, proximity shields, harassment filters, session-time prompts, sickness-aware locomotion defaults) to earn the **EU Trust Label**. This directly targets the experimental pathways that cut sickness and lifted trust.
- (iv). **Institutionalise auditability.** Require **machine-readable safety manifests** and **audit logs** (for moderation/AI decisions) as part of label compliance; align with AI-Act-style audits. This supports oversight **without** undermining interoperability or openness.
- (v). **Run cross-border sandboxes.** Use OPENVERSE to convene **multi-platform pilots** that combine **labels + VCs + audit trails + user-agency** features—using our factorial design as a template for outcome evaluation (trust, openness, comfort, presence). Feed lessons into standards and procurement criteria.

5 DISCUSSION

This chapter interprets the experimental signals through the lens of OPENVERSE's aims and delivery needs. Rather than treating the pilot as a one-off benchmark, we read its effects such as higher trust/privacy and lower sickness under externally signalled rules, no openness penalty, and reduced presence under visible identity as guidance for what to test next, how to fold evidence back into foresight and roadmapping, and which levers are mature enough to operationalise through standards and procurement. Section 5.1 sets out a programme for further research that scales the design across platforms and user groups, strengthens instruments and causal analysis, and packages outputs so they are reusable by the consortium and standards communities. Section 5.2 turns results into scenario-relevant insights for T2.4, showing how gentle, legible safeguards can coexist with creative openness, why accountability should rely on pseudonymous verification rather than blanket identity exposure, and how interoperable safety baselines can travel across worlds. Section 5.3 then converts these insights into policy-ready actions with named owners, active tasks, and concrete means of verification, so that T2.5, WP6, and WP5 can move from evidence to uptake without losing sight of the user experience at the core of an open, human-centric European virtual-worlds ecosystem.

5.1 FURTHER RESEARCH

Building on the evidence from the exploratory study, further research in the field should (i) deepen the experimental programme, (ii) tighten the feedback loops with foresight and roadmapping, and (iii) prepare standardisation and uptake-ready outputs. (iv) Datasets and instruments will be made available via the OPENVERSE Observatory following FAIR principles to support Open Science and replicability. Several contributions to further research activities can be derived from the insight of this deliverable that are directly related to further research activities of the other OPENVERSE WPs and provide insight to tailor further research initiatives, such as:

A. Extend and harden the experimental line (WP3 → WP2/6). Scale the T3.2 exploratory study into a multi-platform, fully crossed design (rules × identity × UX factors), with larger and more diverse samples, and repeated exposures to test persistence of effects. Results should continue to benchmark real platforms against the foresight variables introduced for T2.4, feeding scenario validation and policy options in T2.5. This keeps the experiment anchored to the DoA intention of using live environments to identify behavioural, privacy and identity risks and to compare them with OPENVERSE scenarios.

B. Close the foresight–evidence loop (WP2). As T2.1 to T2.3 deliver ecosystem mapping, marketplace insights and priority areas, their findings, together with refreshed T3.1 trend work, should be reframed as testable hypotheses for the next T3.2 waves and for the sectoral foresight workshops in T2.4. Workshop outputs then flow back into T2.5 to shape the concrete policy/research roadmap (owners, milestones, KPIs).

C. Ethics, legal and IPR co-design (WP3 ↔ WP2/6). Future studies should be co-reviewed with T3.4 (ethics/legal) and T3.5 (IPR/governance) so that instrument refinements (e.g., identity cues, auditability notices) align with forthcoming governance models and with the policy recommendations pipeline in T2.5. The amendment also foresees an ethics advisor and continuous ethics compliance throughout implementation.

D. Standardisation-ready packages (WP6). Translate stable experimental effects into inputs for D6.1 standards gap analysis (e.g., safety-UX baselines, trust labelling, audit trails) and into D6.2 exploitation options. Concurrently, structure research artefacts (safety manifests, identity-credential profiles) so

they can be lifted by relevant SDOs/fora (ISO/IEC SEG on virtual worlds, ITU-T, W3C/IETF tracks) as described in T6.1.

E. Observatory & open evidence (WP6 → WP4/5). Publish datasets, instruments and comparative platform notes into the OPENVERSE Observatory (D6.3/D6.4) to support reuse by partners and by the wider community, while WP4/5 use these materials in dialogues, clustering and joint events.

F. Cross-initiative sandboxes and synergies (WP5). Package the next experiments as “policy sandboxes” combining labels + verifiable credentials + audit trails + user-agency controls, and test them with EU initiatives mapped under WP5 (and NGI communities) to accelerate alignment and uptake across programmes and standards pipelines.

G. Open Science and FAIR Data Availability. In line with Horizon Europe’s Open Science principles and the OPENVERSE Data Management Plan, all anonymised datasets, codebooks, and instruments derived from this study will be curated following the FAIR (Findable, Accessible, Interoperable, Reusable) data principles. Upon validation, these materials will be deposited in the OPENVERSE Observatory (Deliverable D6.3), ensuring long-term accessibility for researchers, policymakers, and standards bodies. Metadata and documentation will be published under an open licence to enable transparent replication, secondary analysis, and integration into future foresight and policy evaluation activities across the consortium.

5.2 USE OF RESULTS OF THE EXPERIMENT IN FUTURE OPENVERSE ACTIVITIES

The experimental insights from T3.2 are best used not as a benchmark against pre-existing scenarios but as formative inputs to the scenario design itself. What follows is a discursive synthesis of the T3.2 findings translated into scenario-relevant implications. The aim is to give facilitators and participants clear, evidence-informed levers to probe when constructing and stress-testing futures for open, human-centric virtual worlds. Throughout, the focus remains squarely on user experience, identity, privacy, openness, and comfort.

5.2.1 From “wild west” spaces to gently structured environments

In the controlled comparisons that juxtaposed a more laissez-faire environment (akin to VRChat) with a more structured one (akin to Horizon Worlds), participants consistently reported higher levels of perceived trust and privacy in the structured setting, together with a marked reduction in simulator sickness. Importantly, this uplift in safety and comfort did not come with any observed penalty to perceived openness or enjoyment during short exposures. Taken together, these effects suggest that users are acutely responsive to visible signals of order, predictability, and care: signage that clarifies acceptable conduct, onboarding cues that prime respectful interactions, and the general ambience of an environment that looks tended rather than chaotic. For scenario design, this indicates a promising direction: futures that feature light-touch but legible safeguards are more likely to generate immediate feelings of safety and inclusivity without dulling creative energy. The practical question for T2.4 is therefore not whether to include structure, but what the minimum viable structure looks like what cues, tutorials, and ambient affordances are sufficient to move the needle on trust and comfort while keeping friction low. A second, subtler implication concerns accessibility: reduced simulator sickness in the structured variant points to a pathway for broader participation, as people who might otherwise churn due to discomfort could remain engaged when comfort is proactively designed in.

5.2.2 Identity visibility versus pseudonymous accountability

The identity manipulation yielded a different pattern. Making identity more visible through verified real-name cues did not increase perceived trust or privacy; instead, it reduced spatial presence, the

felt sense of “being there.” In other words, the promise of visible identity as a shortcut to trust did not materialize in the short-run interactions we observed, whereas the experiential cost was noticeable. This finding does not argue against accountability; rather, it invites a reframing toward pseudonymous accountability. In such a model, users interact through chosen avatars and nicknames, preserving expressive freedom and immersion, while the system maintains verifiable, portable attestations in the background (for example, proof of uniqueness, age eligibility, or community reputation). For scenario construction, this opens a productive middle ground between strict real-name regimes and total anonymity. Scenarios can be articulated around design choices such as which attributes must be verified to sustain trust, when (if ever) identity should become visible to peers, and how due-process safeguards are applied when misconduct occurs. It is also worth exploring domain-specific calibrations: education, healthcare, cultural heritage, and entertainment communities may legitimately select different points on the visibility–privacy spectrum while all relying on the same underlying attestation and redress mechanisms.

5.2.3 Openness and interoperability, tempered by safety baselines

Short exposure studies like T3.2 show that users’ trust signals hinge more on social and safety cues than on how technically open a platform is. That said, greater openness to user-generated content often correlates with higher unpredictability. The policy and design challenge is therefore not to trade openness for safety, but to co-elevate both. Within T2.4, this suggests scenarios that pair openness and interoperability with portable safety primitives: onboarding rituals that travel with the user, consistent reporting and blocking tools, and reputation or attestation signals that can be honored across worlds without exposing personal data. Such patterns allow a creator-friendly ecosystem to flourish while maintaining a dependable baseline of user protections. A useful stress test for scenarios is to ask whether an ‘open’ variant still delivers rapid, legible recourse for ordinary users when something goes wrong; if it does, openness and safety can rise together rather than pull apart.

5.2.4 Synthesis and implications for T2.4 facilitation

Across these axes, three early messages emerge for the foresight team. First, visible safety structure appears to be a powerful and inclusive lever: it boosts perceived trust and reduces physical discomfort without dampening the sense of openness in short sessions. Second, identity visibility is not a reliable proxy for trust and may erode immersion; accountability mechanisms should therefore emphasize verifiable attributes and due-process-ready redress over blanket real-name display. Third, openness and interoperability benefit from portable norms and safety tooling lightweight, platform-agnostic patterns that follow users and creators wherever they go. In workshops, facilitators might therefore prefer scenario variants that combine (a) minimal but legible safeguards, (b) pseudonymous verification with selective disclosure, and (c) safety baselines that are interoperable by design.

To translate this into practical facilitation, consider structuring the scenario canvases around a handful of testable levers. What is the minimal constellation of signals (onboarding, signage, escalation paths) that reliably lifts trust and comfort? Which specific credentials (age, uniqueness, prior conduct) must be verifiable to deter misconduct while protecting privacy, and in which communities do thresholds legitimately differ? What trust signals and user controls can travel across worlds so that ‘open’ ecosystems do not force users to relearn safety from scratch each time? Finally, where might friction or chilling effects begin—at what point do additional requirements start constraining spontaneity, creativity, or inclusion? Answering these questions will help the consortium assemble scenario portfolios that are ambitious yet implementable, and that keep the user’s lived experience at the centre.

5.2.5 Closing note

The overarching value of the T3.2 evidence for T2.4 is pragmatic: it narrows the field of plausible choices by indicating which combinations are most likely to win user trust without compromising immersion and openness. The forthcoming scenarios can therefore start from a position of strength, prototyping futures built around gentle structure, pseudonymous accountability, and interoperable safety, then pressure-testing those futures with stakeholders who bring diverse contexts, constraints, and ambitions. In doing so, the consortium maximizes the chances that the eventual roadmap speaks to what users actually feel and need inside virtual worlds.

After narrowing the field of plausible choices, the T3.2 evidence gives us a practical template for action: combine light-touch, legible safeguards with privacy-preserving accountability and safety primitives that travel across worlds. What follows therefore moves from “what works” in user terms to “what gets done” in project terms. Each recommendation below is framed to be directly liftable into the roadmap, with named owners and active Tasks/WPs, and with means of verification that correspond to artefacts already foreseen in the Description of Action (e.g., the T2.5 roadmap entry and annexes, T6.1 inputs to industry standards, T6.2 exploitation and procurement materials, and the D6.3/D6.4 Observatory pages where registries, conformance reports, and pilot summaries are published). In short, Section 5.3 operationalises the three signals distilled in 5.2.5 (i.e. gentle structure, pseudonymous accountability, interoperable safety), so they can be implemented, inspected, and iterated within OPENVERSE, and reused by policymakers and industry.

5.3 POLICY RECOMMENDATIONS

This section turns OPENVERSE’s priorities into a focused implementation agenda. It distils evidence from the D3.4 experiment and the D2.3 foresight track into concrete measures that can be executed now by the right **owners**, through active **Tasks/WPs**, and against measurable **Means of Verification**. The organising principle is simple: scale what demonstrably lifts trust and comfort **without** narrowing openness, and hard-wire it into standards, procurement, and cross-border pilots. Each recommendation below follows the same cadence, policy rationale in plain language, then a compact list of **Owners / OPENVERSE Tasks / Means of Verification**, so that T2.5 can lift them directly into the roadmap and WP6/WP5/WP4 can drive uptake, transparency, and international alignment. Recommendations from the experiments are:

1) Make “sovereign openness” tangible through visible safeguards (trust label, rule legibility, redress). The priority is to move from experimental insight to a recognisable, certifiable signal that a world is safe, fair, and human-centric without sacrificing openness. A lightweight **EU Trust Label** anchored in pre-entry rule summaries, user rights and redress, and a minimum safety-UX baseline translates that balance into something users see before they commit. Because earlier exploration tasks have finished, content and criteria should now be consolidated through **T2.5** and channelled to **WP6** for standardisation, while **WP4** prepares market-facing guidance. Platforms adopt the label voluntarily at first, but procurement and funding hooks accelerate uptake.

- **Relevant Stakeholders:** EC/DG CNECT (scheme design); CEN/CENELEC/ISO (criteria); Platforms (implementation)
- **OPENVERSE Tasks:** **T2.5** (define label option & KPIs); **T6.1** (translate into specs) ; **T6.3** (project observatory); **WP4** (guidance/briefs)
- **Means of verification:** **T6.3 Project Observatory** (access to the Observatory contents; verifiable downloads).

2) Replace blanket real-name exposure with privacy-preserving credentials (VC/DID). Identity exposure proved a poor lever for trust and a drag on immersion. Accountability should be delivered through **selective-disclosure credentials** (e.g., age/eligibility/uniqueness) compatible with eIDAS-aligned wallets, not through visible real names. The narrative for policy, standards, and procurement is simple: meet the assurance need while keeping presence high. Delivery shifts to **T2.5** to specify the policy option and pilot shape; **WP5** Collaboration and Synergies and **T6.1** Input to industry standards.

- **Relevant Stakeholders:** EC/eIDAS2 team; Member State authorities (recognition); Platforms (integration)
- **OPENVERSE Tasks:** **T2.5** (policy option & pilot framing); **T6.1** (credential profiles/interoperability notes); **T6.2** (exploitation/procurement hooks)
- **Means of verification:** **Pilot evaluation file set (T2.5)** containing anonymised **verification logs** exported to **T6.3** Project observatory.

3) Standardise a safety-by-design UX baseline. Trust improves when comfort improves; comfort improves when safety tools are predictable and easy. Codify a **reference safety-UX** (panic/exit, proximity shields, harassment filters, session-time prompts, sickness-aware locomotion) with verifiable conformance tests. The baseline becomes the entry ticket to the label and to public calls. While earlier UX experimentation informed the content, the operational work now sits with **WP6** and **T2.5** to fix the baseline and link it to adoption levers.

- **Relevant Stakeholders:** SDOs (CEN/CENELEC/ISO); Platforms; EC (policy linkage)
- **OPENVERSE Tasks:** **T6.1** (baseline spec & tests as input to industry standards); **T2.5** (policy recommendations); **T6.2** (exploitation and sustainability actions)
- **Means of verification:** **T6.1 Input to Industry standards** (through relevant initiatives undertaken in the task) archived and displayed in **T6.3** (Project Observatory).

4) Institutionalise auditability (safety manifests; moderation/AI audit logs and explainability). Users and regulators need to see what safety exists and how it operates. Increase transparency via a **machine-readable Safety Manifest** and **tamper-evident audit logs** for significant moderation/AI actions, plus plain-language “why” notices. Manifests gate label eligibility and create a direct evidence channel to oversight and the Observatory. With research phases closed, operationalisation rests with **WP6** (schemas/standards), **T2.5** (policy requirement), and **WP4** (guidance to regulators and industry through communication and dissemination).

- **Relevant Stakeholders:** EC & national regulators; Platforms
- **OPENVERSE Tasks:** **T6.1** (input to industry standards); **T2.5** (roadmapping); **T6.3** (observatory); **WP4** (implementation guidance through communication and dissemination actions)
- **Means of verification:** **Machine-readable Safety Manifests published in the Observatory (T6.3)** with validator receipts

5) Stand up a pan-EU regulatory sandbox (“28th regime”) to de-risk scaling. Evidence-to-policy is fastest when testing and supervision travel together. A **borderless sandbox** lets platforms (especially SMEs) trial the label, credentials, manifests, and safety-UX under a harmonised charter instead of 27 variants. The consortium’s role is coordination and methodology: **T2.5** frames the recommendations, **WP5** brokers cross-initiative participation through collaboration and synergies, and **WP6** ensures test outputs flow into industrial standards.

- **Relevant Stakeholders:** EC with Member State regulators; Platforms (participants)
- **OPENVERSE Tasks:** **T2.5** (recommendations); **WP5** (collaboration and synergies); **T6.1** (input to industry standards); **T6.3** (publish sandbox results through the observatory)
- **Means of verification:** **WP5 Sandbox Collaboration and Synergies** (relevant synergies activated through WP5 tasks).

6) Convert pilots into standards and procurement (create market pull). What works must become what's bought. Channel stable findings into **pre-normative deliverables** and **model procurement clauses**. This turns the label, credentials, manifests, and safety-UX into requirements that matter in tenders and calls, avoiding voluntary-only drift. With upstream tasks finished, the operational engine is **WP6** (standards and exploitation) working with **T2.5** to keep policy alignment.

- **Relevant Stakeholders:** SDOs; EC and MS procurement bodies
- **OPENVERSE Tasks:** **T6.1** (Input to Industry standards); **T6.2** (Exploitation and sustainability); **T2.5** (implementation track in the roadmap)
- **Means of verification:** **SDO submission receipts & procurement clause repository** (T6.1/T6.2), with document IDs and public links logged in **T6.3** (project observatory).

7) Orchestrate the governance cadence via T2.5. The roadmap must now operate like a programme: clear owners, milestones, KPIs, and review cycles published to the Observatory. **T2.5** leads this cadence, convenes validation workshops, and issues periodic updates so policymakers and industry can see progress, slippage, and course correction. Earlier WP2 mapping and foresight remain inputs; the operational drive is now scheduling, accountability, and public traceability.

- **Relevant Stakeholders:** EC and MSs decision makers and regulators
- **OPENVERSE Tasks:** **T2.5** (roadmapping); **T6.3** (observatory); **WP4** (briefs/events through communication and dissemination)
- **Means of verification:** **Versioned Roadmap Matrix** (PDF/CSV) **published in T6.3** project observatory.

8) Finance what works; protect SME time-to-conformance. SMEs will carry much of the creative load but face the steepest compliance curve. The project should promote adoption with **targeted vouchers** (usability/accessibility testing, security hardening, VC integration) and **platform credits/SDKs** so meeting label criteria is fast and affordable. The aim is not subsidy for its own sake, but **time-to-conformance** reduction tied to measurable user-level outcomes.

- **Relevant Stakeholders:** EC/MS innovation agencies; EIB/EIF; Platforms
- **OPENVERSE Tasks:** **T6.2** (exploitation & financing pathways); **T2.5** (recommendations); **WP4** (SME outreach through communication and dissemination)
- **Means of verification:** **T6.2 Exploitation and sustainability** (credible exploitation financing pathways with NdA signed with potential investors).

9) Align internationally to avoid fragmentation. Virtual worlds are global; safety and identity should travel across borders. **Mutual recognition** of label criteria and credential profiles, and on publishing **open test methods** so results are comparable. **WP5** leads the dialogues through collaboration and

synergies; **T6.1** maintains spec alignment with ISO/IEC/IEEE/ITU/W3C through input to industry standards; **T2.5** maps equivalence in the roadmap to keep Europe open yet confident.

- **Relevant Stakeholders:** EC (trade/standardisation); Member States; SDOs; Platforms
- **OPENVERSE Tasks:** **WP5** (collaboration and synergies); **T6.1** (spec alignment for industry standards); **T2.5** (roadmapping)
- **Means of verification:** **Signed mutual-recognition MoUs / liaison statements (WP5)** archived and indexed in **T6.3** (observatory), with SDO work-item numbers where applicable.

6 CONCLUSIONS

The exploratory study shows that modest, well-signalled governance raises perceived trust and privacy while reducing simulator sickness, without depressing perceptions of openness. Across 101 participants, trust/privacy increased from pre- to post-exposure irrespective of condition, indicating an onboarding/education effect from short, structured sessions. More importantly for policy design, externally signalled rule frameworks consistently outperformed laissez-faire settings on trust/privacy and comfort and did so without any measurable openness penalty such as an empirical manifestation of the “sovereign openness” idea developed in OPENVERSE foresight work. These results suggest that visible safeguards and rule legibility are high-yield levers for user experience: when comfort improves, trust rises; when environments look tended and predictable, users feel safer even as creative latitude remains intact.

Identity policies produced a more selective pattern. Requiring visible, verified identity did not raise trust or openness and reliably reduced spatial presence, whereas pseudonymity preserved immersion without observable safety losses in this setting. Taken together with the rule-framework effects, this indicates that accountability should be engineered via privacy-preserving assurance mechanisms (e.g., selective disclosures and background attestations, rather than blanket real-name display). In short exposures, users rewarded clarity about rules and recourse more than they rewarded identity visibility, and the correlational structure reinforced comfort as a hinge variable linking sickness, presence, and trust.

Methodologically, the study demonstrates the feasibility of translating foresight axes into testable manipulations and interpretable effect sizes. The factorial design, balanced cells, and conservative data-quality procedures yielded stable primary composites (trust/privacy, openness; SSQ) while flagging developmental instruments (presence; baseline privacy) for refinement. Limitations related to pilot sampling, brief exposures, and scale reliability issues on two constructs, limit strong generalisations but do not undercut the central signals on rule visibility, comfort, and identity policy. As such, the evidence is fit for purpose as an input to roadmapping: it narrows the choice set toward visible safeguards, comfort-first safety UX, and privacy-preserving credentials, and away from identity exposure as a proxy for trust.

The results justify advancing an EU-facing trust label anchored in pre-entry rule summaries, user rights and redress, and a safety-UX baseline; codifying auditability through machine-readable safety manifests and explainable moderation trails; and piloting privacy-preserving credentials that meet assurance needs without eroding immersion. They also support a pan-EU sandbox approach to de-risk scaling and a standards-and-procurement pipeline to convert what works into market pull. In this way, T3.2 evidence becomes a bridge from scenario to specification, from specification to recommendations, and from recommendations to uptake, keeping the user’s lived experience at the centre while delivering concrete levers for T2.5, WP6, and WP5 to execute.

REFERENCES

European Commission. (2023). *An EU initiative on Web 4.0 and virtual worlds: A head start in the next technological transition (COM(2023) 442)*. Brussels.

European Commission, Digital Strategy. (2023/2025). *Virtual Worlds - Digital Identity; Virtual Worlds fit for people*. Brussels.

Laiz, H., et al. (2024). *The virtual worlds: Privacy and information security risks*. International Journal of Information Management Data Insights, 5(2), 100373.

OPENVERSE. (2025). *Deliverable D2.3, Report on technological foresight*.

OPENVERSE. (2024). *Deliverable D3.1, Report on Virtual Worlds Trends and Benchmarking*.

Polychronaki, M., et al. (2024). *Decentralized identity management for virtual-worlds-enhanced education: A literature review*. Electronics, 13(19).

Veltri, G. A., Mureddu, F., Innocenti, A., & Sirizzotti, M. (2025). *Regulatory frameworks and digital identity in the virtual worlds: An exploratory VR experiment on trust, privacy, and openness* (Draft findings under preparation for peer reviewed scientific journal).

VRChat. (2022). *Safety and Trust System (Trust Rank & user safety features)*.

World Economic Forum. (2023). *Privacy and safety in virtual worlds* (Insight Report).

Wu, H. (2023). *Digital identity, privacy security, and their legal safeguards in the virtual worlds*. Security and Safety, 2(1).

XR Safety Initiative (XRSI). (2020–2022). *XR Privacy Framework; Virtual Worlds Safety Week outcomes*.

ANNEX I: THE TWO VIRTUAL WORDS USED IN THE EXPERIMENT

Meta Horizon Worlds

Horizon Worlds is Meta's social, user-generated "metaverse" experience. People create and explore worlds; attend events, games, and hangouts; and socialize with avatars.

- Cross-platform access: Available on Quest headsets and expanding to mobile and web in supported regions.
- Creation tools: In-world building tools let you prototype spaces, game mechanics, and interactive experiences without leaving VR.
- Communities & events: From casual meetups to live shows and game nights; moderators, reporting, and safety tools help keep interactions positive.
- Good for: Team social events, audience engagement, low-barrier prototyping of concepts and exhibits, and broad community outreach.

VRChat

VRChat is a long-running social VR platform centered on deep customization and community-built content.

- Platforms: Works on Quest (standalone), PC-VR (SteamVR), and desktop (no headset required), which helps broaden participation.
- Creation stack: Powerful pipelines (Unity-based worlds and avatars) and Udon visual scripting allow complex interactions, mini-games, and role-play systems.
- Safety & trust: Robust muting, blocking, and safety layers with granular control over audio/visual effects, avatars, and interactions.
- Good for: Highly customized worlds, avatar-driven communities, meetups that mix desktop and VR users, and experimental, creator-led events.




Figure 3 Screenshot from the experiment

ANNEX II: THE HEADSET USED IN THE EXPERIMENTS

Meta Quest is a standalone XR headset designed to deliver virtual reality (VR) and mixed reality (MR) without tethers, base stations, or external trackers. Its inside-out tracking uses onboard cameras to map your environment and track movement, while ergonomic touch controllers provide precise input with haptics. You can also use controller-free hand tracking for casual navigation, media, and select apps—useful when you want quick interactions or a more natural feel.

The headset’s defining capability is mixed reality through full-color passthrough. Instead of isolating you in VR, Quest blends digital content into your actual room so you can see walls, furniture, and people while placing 3D windows, tools, or game elements around you. This unlocks room-scale gameplay (think tabletop strategy on your coffee table or fitness apps that use your real space) and practical productivity scenarios like pinning multiple virtual screens over your desk. Boundary tools help you quickly set safe play areas, and spatial anchors let apps remember where you placed things in your room.

Under the hood, Quest 3 uses Qualcomm’s Snapdragon XR2 Gen 2 platform for a substantial step up in graphics, loading times, and scene understanding compared to prior generations. Recent software updates have doubled down on everyday usability: a more flexible home/desktop layout, improved windowing and multitasking (pin a browser next to a notes app or email), and tighter PC-VR integration. With Meta Link (USB-C) or Air Link (Wi-Fi), you can stream high-fidelity PC VR titles from your desktop while still enjoying the simplicity of standalone apps when you’re away from your desk.

The Quest ecosystem is mature and diverse. The Quest Store offers a wide library of games—from rhythm workouts and narrative adventures to realistic sims—alongside a growing set of productivity tools (virtual monitors, whiteboarding, design reviews), fitness titles (guided programs, heart-pumping workouts, mixed-reality boxing), and creative apps (sculpting, spatial notes, 3D sketching). On the social side, platforms like Meta’s Horizon Worlds and VRChat enable meetups, events, collaborative building, and a thriving creator culture. Whether you’re co-working in a virtual studio, attending a live talk, or hopping into a casual hangout, Quest makes it easy to discover communities and content.

Day-to-day, Quest aims to be fast to start and simple to live with: quick guardian setup, easy account sign-in, cloud backups for your library, and parental/guardian tools where needed. Developers benefit from robust SDKs for hand tracking, passthrough, spatial anchors, and mixed-reality compositing—so new apps continue to feel smarter and more “room-aware” over time. For power users and teams, the headset’s blend of standalone convenience plus PC-VR horsepower covers both casual sessions and demanding use cases like sim racing or design review.

In short: Meta Quest delivers an accessible entry point to spatial computing—no wires, no base stations, and no fuss—while increasingly supporting serious productivity, creation, and social presence. If you want one device that plays great standalone games, supports mixed-reality workflows, and can tap into PC-grade VR when you need it, Quest is the most well-rounded option in its class.

ANNEX III: EXPERIMENTS' QUESTIONNAIRES AND EVALUATION SCALES

Pre-Experiment Measurements

Demographic Data

- Age (open field).
- Gender (Male, Female, Non-binary, Prefer not to answer, Other).

Questionnaire on Familiarity with Virtual Reality (1 = Strongly Disagree ... 7 = Strongly Agree)

- VR1. I regularly use VR headsets to explore immersive environments.
- VR2. I am confident in customizing avatars or digital identities in virtual worlds.
- VR3. I frequently participate in multi-user virtual platforms (e.g., Horizon Worlds, Roblox, Second Life).
- VR4. I am familiar with the privacy and security settings of VR applications.
- VR5. I keep myself updated on news about VR technology and its regulations.
- VR6. I can quickly adapt to different controllers or modes of VR interaction.

1.3 Basic Attitudes on Privacy and Trust (1 = Strongly Disagree ... 7 = Strongly Agree)

- P1. I am concerned about how commercial platforms use my personal data.
- P2. I trust government regulations more than platform policies to protect my data in VR.
- P3. I prefer to remain anonymous when interacting online.
- P4. Verifying my real identity in digital spaces makes me feel safer.
- P5. I usually read the terms and conditions before using new VR services.
- P6. I am willing to share personal data if the regulatory framework is clear and enforced.

Post-Exploration Questionnaires

Perceptions of Trust and Privacy (1 = Strongly Disagree ... 7 = Strongly Agree)

- TP1. I felt safe while interacting in the VR environment.
- TP2. I trusted the regulatory body (platform rules or government regulation) to protect my privacy.
- TP3. Identity verification cues (e.g., verified badge) increased my sense of security.
- TP4. I was concerned that my actions might be recorded without my consent.
- TP5. I would use this environment again with my real account.
- TP6. I would recommend this VR experience to friends who care about privacy.

Aspects of Openness and Free Access Innovation (1 = Strongly Disagree ... 7 = Strongly Agree)

- OS1. I was able to access public spaces and talk to others without additional costs or the obligation to join exclusive groups.
- OS2. It was easy to discover new people and join their conversations or activities.
- OS3. I was able to invite friends or colleagues into the same VR space without complicated steps.
- OS4. The platform made it easy to share objects or experiences I created so that others could use or modify them.
- OS5. I had the impression that ordinary users, not just the company, can influence future features and community rules.
- OS6. I would like to organize my own events or share creations with this VR community in the future.

2.3 IPQ (1 = Strongly Disagree ... 7 = Strongly Agree)

IPQ1. How aware were you of the real world around you while navigating the virtual world (e.g., sounds, room temperature, other people, etc.)?

IPQ2. How real did the virtual world seem to you?

IPQ3. I had the feeling of acting within the virtual space rather than from outside.

IPQ4. How consistent did your experience in the virtual environment feel with your real-world experience?

IPQ5. How real did the virtual world seem to you?

IPQ6. I did not feel present in the virtual space.

IPQ7. I was not aware of my real environment.

IPQ8. In the computer-generated world, I had the feeling of “being there.”

IPQ9. I somehow felt that the virtual world surrounded me.

IPQ10. I felt present in the virtual space.

IPQ11. I still paid attention to the real environment.

IPQ12. The virtual world seemed more realistic than the real world.

IPQ13. It seemed to me that I was only perceiving images.

IPQ14. I was completely captivated by the virtual world.

SSQ (1 = None ... 4 = Strong)

Simulator Sickness Questionnaire (SSQ)

Symptom	1 None	2 Slight	3 Moderate	4 Strong
Nausea				
Dizziness				
Headache				
Fatigue				
Sweating				
Disorientation				
Eye strain				
Difficulty concentrating				
Increased salivation				
Blurred vision				
Double vision				
General discomfort				

ANNEX IV: ETHICS ASSESSMENT BY ETHICS ADVISOR

6.1 ETHICS ADVISOR ASSESSMENT REPORT – DELIVERABLE D3.4

Project: OPENVERSE (Grant Agreement No. 101135701)

Deliverable: D3.4 – Report on Virtual Worlds Environments Exploratory Study

Evaluator: Ethics Advisor (Independent)

Date: 14/10/2025

6.1.1 Overview

This report presents the ethics advisor's assessment of Deliverable D3.4 in relation to the Horizon Europe Ethics Guidelines and the OPENVERSE Ethics Framework. The evaluation focuses on data protection, participant safety, informed consent, equality, inclusivity, and adherence to the EU Charter of Fundamental Rights.

6.1.2 Summary of Ethical Compliance

D3.4 demonstrates strong compliance with EU ethical principles and OPENVERSE's internal ethics framework. The exploratory study integrates ethical safeguards across participant engagement, data collection, and analysis.

Ethical Area	Assessment	Evidence / Comments
Data Protection (GDPR)	Compliant	Aggregate data only, pseudonymised, no identifiers stored.
Informed Consent	Compliant	Plain-language consent forms, withdrawal rights guaranteed.
Psychological Safety	Compliant	Simulator Sickness Questionnaire and opt-out mechanism used.
Equality & Inclusion	Strong	Gender-balanced sample; equity metrics used for comfort.
Transparency & Accountability	High	Methodological documentation and annexes open for review.

6.1.3 Sensible Points and Mitigation Measures

Area of Sensitivity	Ethical Concern	Observation in D3.4	Mitigation / Recommendation
Data logs	Behavioural data may contain indirect identifiers	Logs aggregated and anonymised	Include public DPIA summary (D7.7 alignment).
Immersive exposure	Simulator discomfort risk	Monitored with SSQ and acclimatisation	Maintain facilitator supervision record.
Identity exposure	Avatar/real ID linkage risk	Handled as ethical variable	Maintain pseudonymisation and Ethics Board review.
Accessibility	Limited representation of users with disabilities	Focus on comfort equity	Expand inclusion to users with sensory/motor limitations.

6.1.4 Overall Ethical Acceptability

D3.4 meets Horizon Europe ethical requirements and demonstrates operational adherence to the Ethics Guide for Advisors. The deliverable applies 'ethics-by-design' principles and integrates monitoring measures ensuring participant safety and legal compliance.

Criterion	Evaluation
Adherence to EU Ethical Requirements	High (Compliant)
Transparency & Accountability	High
Participant Welfare & Safety	High
Inclusivity & Equality	Strong, minor improvements possible
Residual Risk Level	Moderate (identity exposure, fatigue)
Final Ethical Verdict	Acceptable – Green/Green

6.1.5 Recommendations for Future Actions

1. Ensure formal Ethics Advisor sign-off included in Annex IV.
2. Reinforce psychological-safety measures during VR sessions.
3. Expand accessibility tests to include participants with disabilities.
4. Maintain the Ethics Issue Register under WP7 Quality Assurance Plan.

6.1.6 Conclusion

Deliverable D3.4 is ethically acceptable and consistent with the Horizon Europe framework. No major risks have been identified. Minor corrective actions, if implemented, will ensure full compliance and elevate the deliverable to a best-practice reference in human-centric XR experimentation.

ANNEX V: DPIA SUMMARY

6.2 DATA PROTECTION IMPACT ASSESSMENT (DPIA) SUMMARY – DELIVERABLE D3.4

Project: OPENVERSE (Grant Agreement No. 101135701)

Deliverable: D3.4 – Report on Virtual Worlds Environments Exploratory Study

Prepared by: Ethics Advisor

Date: 14/10/2025

6.2.1 Purpose of the DPIA

The Data Protection Impact Assessment (DPIA) for D3.4 evaluates the potential privacy risks arising from the exploratory study of virtual world environments conducted under WP3, Task 3.4. The DPIA ensures compliance with Regulation (EU) 2016/679 (GDPR), particularly Articles 5, 6, 9, and 35, and aligns with the CNIL methodology and the European Data Protection Board (EDPB) guidelines (WP248 rev.01). The study involved limited personal data processing from participants interacting in controlled XR/VR environments.

6.2.2 Description of the Processing

Element	Description
Processing Activities	Collection of user feedback, comfort metrics, and interaction data during VR test sessions.
Data Subjects	Adult volunteers (18+) participating in short immersive trials (5–10 minutes).
Data Types	Non-identifiable interaction data, anonymised questionnaire responses, session logs without personal identifiers.
Purpose	Evaluation of usability, ethical perception, and accessibility of virtual environments.
Lawful Basis (Art. 6 GDPR)	Explicit informed consent obtained prior to participation.
Special Categories (Art. 9 GDPR)	None processed; physiological or biometric data were not recorded.

6.2.3 Risk Assessment

Risk Area	Potential Impact	Likelihood	Mitigation Measures
Re-identification of participants	Minimal due to pseudonymised data but possible via behaviour traces	Low	Data aggregated and anonymised; no audio/video stored.
Uninformed consent or misunderstanding	Possible minor risk of under-informed participation	Low	Plain-language consent forms and opt-out rights ensured.

Psychological discomfort	Transient simulator sickness or anxiety	Medium	Pre-screening, session limit, and post-experience debrief implemented.
Data misuse or unauthorised access	Low risk due to limited dataset	Low	Data stored on EU servers with restricted access; deleted after analysis.

6.2.4 Data Flow and Storage

Data were collected locally during test sessions, stored on encrypted devices, and transferred securely to a GDPR-compliant EU cloud repository. Anonymised datasets were analysed by the University of Galway (INS) and deleted after completion of the study. No data were shared with external partners or transferred outside the EU/EEA.

6.2.5 Data Retention and Deletion

Data are retained only for the duration required for analysis and reporting (maximum 12 months). Raw logs and session-level identifiers are permanently deleted after aggregation. Derived, anonymised datasets may be archived for transparency and reproducibility in accordance with the FAIR principles.

6.2.6 Consultation and Approval

The DPIA has been reviewed and validated by the project's independent Ethics Advisor and the Data Protection Officer of the coordinating institution. Residual risk is considered low. A summary of this DPIA will be published in OPENVERSE DMP and referenced in the project's public Data Management Plan.

6.2.7 Conclusion

The DPIA confirms that the processing activities in D3.4 pose a low risk to the rights and freedoms of participants. All data were processed under informed consent, using anonymisation and secure storage. No sensitive personal data were collected, and all procedures comply with GDPR and Horizon Europe ethical standards.