CoSTAR National Lab · Business Futures

Performance, Participation, Provenance and Reward (P3R) in Responsible AI

Exploring Governance in AI-Enabled Live Event Contexts

Findings from WP2 Focus Group Research

Version: **1.1**
Date: **May 2026**
Authors: **P3R WP2 Research Team**
Status: **Final**
Programme: **CoSTAR Business Futures · P3R**

Contents

1. Executive Summary 3
2. Introduction and Context 4
2.1 Research Design 4
3. Findings 5
3.1 What People Value 5
3.2 The Baseline Consent Architecture 6
3.3 What Governance Unlocks – and Where It Fails 10
3.4 Fairness and Compensation 13
3.5 Identity Boundaries 14
3.6 Configurations and Structural Limits 17
3.7 Location Comparison and Heterogeneity 19
4. Implications for Responsible AI Design 20
4.1 Five Architectural Requirements 20
4.2 Implications for Identity Research 21
Appendix A: Methodology 23
A.1 Research Design 23
A.2 Participants 23
A.3 Exercise Structure 23
A.4 Analytic Approach 23
A.5 Interpretive Constraints 24
Appendix B: Reference Tables 25
Table B.1: Experiential Moments (A–N) 25
Table B.2: Consent/Governance Mechanisms (1–10) 26
Table B.3: Compensation Mechanisms (1–10) 26
Table B.4: Identity/Likeness Use Cases (1–5) 27
Table B.5: Participants Count by Exercise 27

1. Executive Summary

AI-enabled live events promise immersive, personalised experiences – but they depend on collecting and using personal data, including biometric and identity information, in ways that existing governance frameworks were not designed to address. This report presents findings from Work Package 2 (WP2) of the P3R programme, which investigates how people evaluate consent, control, compensation, fairness, and identity in these prospective contexts. Our focus is not on abstract privacy attitudes, but on how people respond to specific AI-enabled experiences when presented with concrete safeguards.

Seventy-three participants attended across two UK locations (England, n = 48; Scotland, n = 25, see Appendix B.5 for details) responded to a Design Fiction video depicting a future AI-enabled live event. A sequential exercise structure – covering value perception, baseline data-sharing willingness, governance impact, compensation preferences, and identity boundaries – generated structured quantitative data supplemented by qualitative discussion. See Appendix A for methodology and Appendix B for the full reference tables.

Six principal findings emerge from the analysis:

- What matters most is hardest to unlock. Only two experiences achieved top-priority status: Accessibility features (66.7%) and Venue check-in (58.0%). Both face high initial resistance and both remain resistant to governance alone. For Venue check-in, the resistance is not to frictionless entry but to the data infrastructure that enables it.
- People are cautious by default. No data type or stakeholder achieves majority consent before governance is introduced. Even the most trusted actor – the venue – falls below 50% (49.7%). A 37.7 percentage-point gap separates organisational actors from the general public.
- Governance helps, but not enough. Safeguards improve willingness in 34% of the 275 data-sharing scenarios tested, but 66% remain resistant. Three mechanisms account for more than half of all changes in willingness: Anonymisation (21.4%), time-limited/revocable consent (19.0%), and verified technical controls (16.0%).
- Fairness means payment – but not universally. Revenue share and direct cash payments dominate compensation preferences (both at 87.7%). However, 12.3% of participants reject monetary compensation entirely, believing their data and likeness should not be treated as something to be bought and sold.
- Data sharing and identity are governed by different logics. Willingness to allow identity/likeness use follows distinct patterns from data sharing and requires separate governance. The apparent paradox in Avatar Representation in Virtual Spaces – avatars partially working yet a hard limit persisting – resolves as two sub-populations: Those who object to real-likeness exposure, and those who object to virtual-space digital presence in any form.
- Some boundaries are structural. 182 of 275 scenarios remain resistant after governance; 1 of 5 identity use cases carries an absolute red line; and findings diverge across both locations in all four research domains. These are not implementation failures but hard limits that Responsible AI design must work within.

2. Introduction and Context

AI technologies are increasingly embedded in live event environments – from wearable devices and avatar creation to real-time crowd analytics and personalised content. While these technologies promise enhanced experiences, they depend on the collection, processing, and use of personal data at a scale and intimacy that existing consent frameworks were not designed to address, and where identity and likeness – not just informational data – may be captured and repurposed. WP2 of the P3R programme investigates these challenges empirically, examining how people respond to specific AI-enabled experiences when presented with concrete safeguards, rather than their general views on privacy.

2.1 Research Design

A Design Fiction video was created as a provocation depicting a future AI-enabled live event. This was presented to focus groups across two UK locations (England and Scotland), followed by a sequential exercise structure comprising: (a) A value perception sprint to identify which experiences matter most; (b) a data portrait mapping baseline willingness to share; (c) a consent shift challenge measuring the impact of governance mechanisms; (d) a compensation ranking capturing fairness expectations; (e) a likeness and image rights exercise mapping identity boundaries; and (f) a free discussion on Responsible AI. All groups followed the same protocol.

The exercises were structured around four building blocks, summarised here and detailed in full in Appendix A:

- 14 experiential moments – key scenes from the Design Fiction covering wearable setup, avatar creation, venue check-in, crowd visuals, friend finding, accessibility features, content licensing, digital marketplace, and more.
- 10 governance mechanisms – ranging from time-limited and revocable consent to anonymisation safeguards, granular opt-in controls, transparent audit trails, and verifiable technical certifications.
- 10 compensation mechanisms – including direct cash payments, revenue share/royalties, platform credits, exclusive access, micro-payments, and charitable donation options.
- 5 identity/likeness use cases – covering promotional material, avatar use in virtual spaces, AI after-movies, social media content, and personal highlight reels.

Quantitative analysis pools all active participants; location comparison identifies where findings diverge. Per-group qualitative analysis is retained for discussion dynamics.

3. Findings

Our findings follow a progressive narrative structure that integrates across the five exercises: What participants value, the starting-point consent landscape, what governance unlocks and where it fails, compensation and fairness expectations, and the distinct domain of identity and likeness. The final two subsections draw these threads together into a classification of how different experiences respond to combinations of interventions and examine location-level stability.

3.1 What People Value

> *'I think this is an incredible opportunity for accessibility for people who like yourself is not a big crowds fan, for people who are like in wheelchairs and that is anxiety inducing, for people internationally who are in countries where there isn't a lot of music, that sort of thing. I think that could be really amazing.' (FG05-P02)*
>
> *'it would literally just be streamlining, you know. It takes so long to line up for a venue and wait, if there's any way to make that quicker, then I feel like that's welcome by anyone.' (FG06-P02)*

Two moments achieved top-priority status across both locations: Accessibility features (66.7%) and Venue check-in (58.0%). These are the only experiences selected by at least half the sample and rated highly important by those who chose them. Both are shared across the two data-collection locations, providing a stable anchor for subsequent analysis. However, 6 of the 14 moments overall are ranked differently between locations, indicating that value priorities are only partially location-stable.

Classification used two metrics: How widely an experience was selected (breadth), and how highly it was ranked by those who selected it (intensity). Both thresholds must be met simultaneously for Tier-1 status – a moment selected by everyone but ranked last would not qualify, nor would one ranked first by a tiny minority.

Not all value is convergent. Most moments show meaningful ranking polarisation among selectors – people who selected the same moment often disagreed sharply about how to rank it. This has a practical implication: Collapsing priorities into a single ranked list obscures both the contested cases and the niche ones. A moment ranked moderately by everyone looks identical in a ranked list to one ranked first by some and last by others; and one experience (Digital marketplace, selected by only 18.8%) would be invisible in a headline ranking despite being an intensely held priority for those who chose it.

Figure 1 maps all 14 moments by breadth of selection (x-axis) and average importance rating among selectors (y-axis); only the two moments in the top-right quadrant meet both Tier-1 criteria simultaneously.

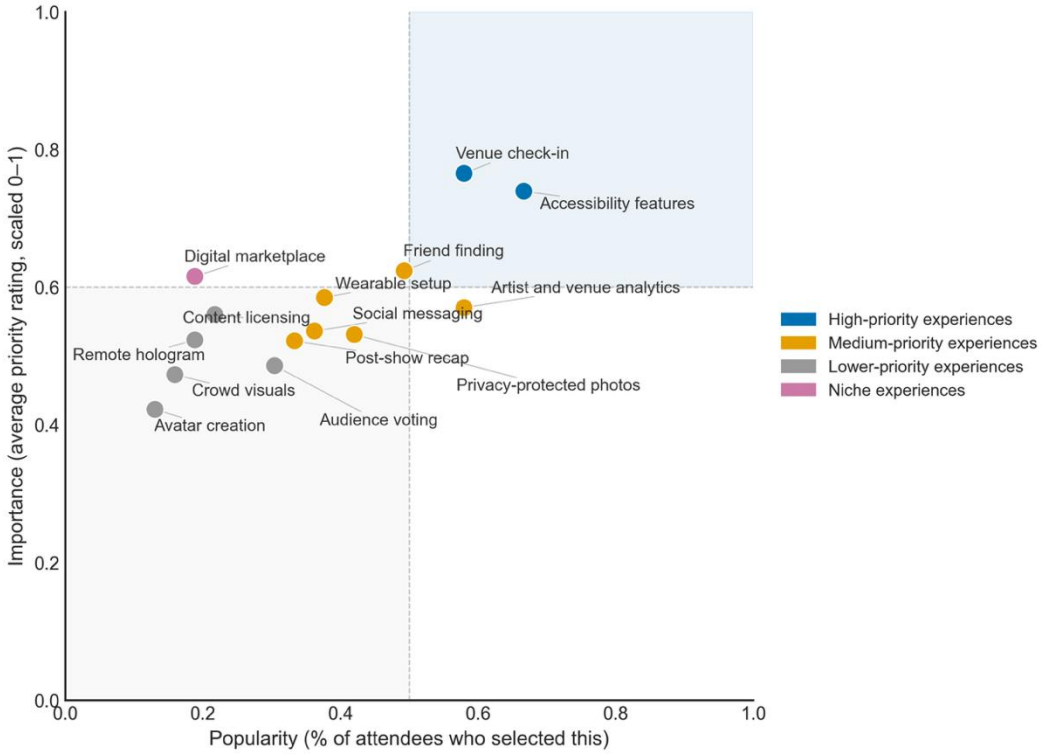

Figure 1. Value Perception (Breath × Intensity).

3.2 The Baseline Consent Architecture

> *'if it's just your username, I feel that's fine… because that's my data, but it's not like personal, private data…And then when it comes to financial data, you are going have to give it to the venue if you want to pay to get in as well. So it's like, there's stuff that you kind of have to, not have to but to a certain extent.' (FG09-P02)*

Before any governance mechanisms were introduced, participants indicated their willingness to share data across a matrix of 14 moments × up to 11 data types × 5 stakeholders, producing 275 individual data-sharing scenarios per participant. A scenario is a specific combination of one moment, one data type, and one stakeholder – the fine-grained unit of the consent matrix. This matrix captures the baseline consent landscape that governance must navigate.

3.2.1 Stakeholder Trust Gradient

Participants are substantially more willing to share data with institutional actors than with diffuse publics, with a 37.7 percentage-point gap from venue (49.7%) to general public (12.0%). Neither endpoint achieves majority consent, but the gradient between them has a clear logic: Institutional actors have defined roles, known purposes, and accountability; diffuse publics have none of these.

The full rank order, from most to least trusted:

- Venue or organiser: 49.7% – the most trusted actor, but still below majority consent.
- Commercial partners: 47.3%.
- Artist or performer: 32.9%.
- Other fans: 28.8%.

- General public: 12.0% – near-universal refusal.

Figure 2 shows pooled baseline willingness by stakeholder and illustrates a gradient that is not a smooth decline. There is a sharp drop between artist/performer and other fans, and again at the general public, reflecting a qualitative difference between institutional and diffuse categories. Both locations agree on this rank order; the gradient is a location-stable feature of the data.

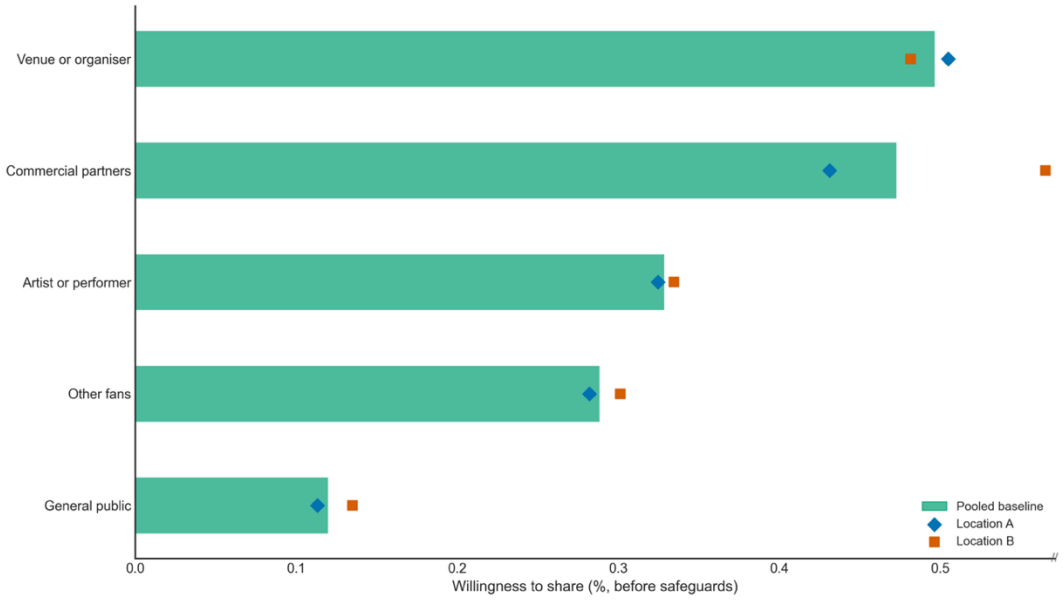

Figure 2. Stakeholder Trust.

The near-parity between commercial partners and the venue is the most counterintuitive feature of this gradient and merits interpretive caution. Workshop discussions during the research process suggest that participants may have read 'commercial partners' through the lens of platform access – understanding consent to these operators as a condition of using the ConsentWave system's features rather than as a genuinely open choice. If so, the 47.3% figure may reflect pragmatic acceptance of a perceived access requirement rather than an expression of trust in commercial partners as a category of data recipient. Willingness to consent under a perceived condition of access is analytically distinct from willingness rooted in genuine trust, and the governance implications of each differ considerably. Future research should probe this distinction directly, for example by varying how commercial partners are framed and whether platform access is presented as contingent on consent to them.

3.2.2 Data-Type Sensitivity

No data type achieves majority consent at baseline. Even the least sensitive category – identity data (name, user ID, avatar profile) at 43.4% – falls short. Biometric data (facial recognition, gait, movement patterns) is among the most refused at 23.0%, and account and login data (passwords, authentication tokens) sits at 18.7% at the foot of the hierarchy. Default caution is universal as not a single type crosses the 50% threshold.

Figure 3 ranks all data types by pooled baseline willingness from least to most sensitive; the 50% majority-consent threshold, which no type reaches, is marked for reference.

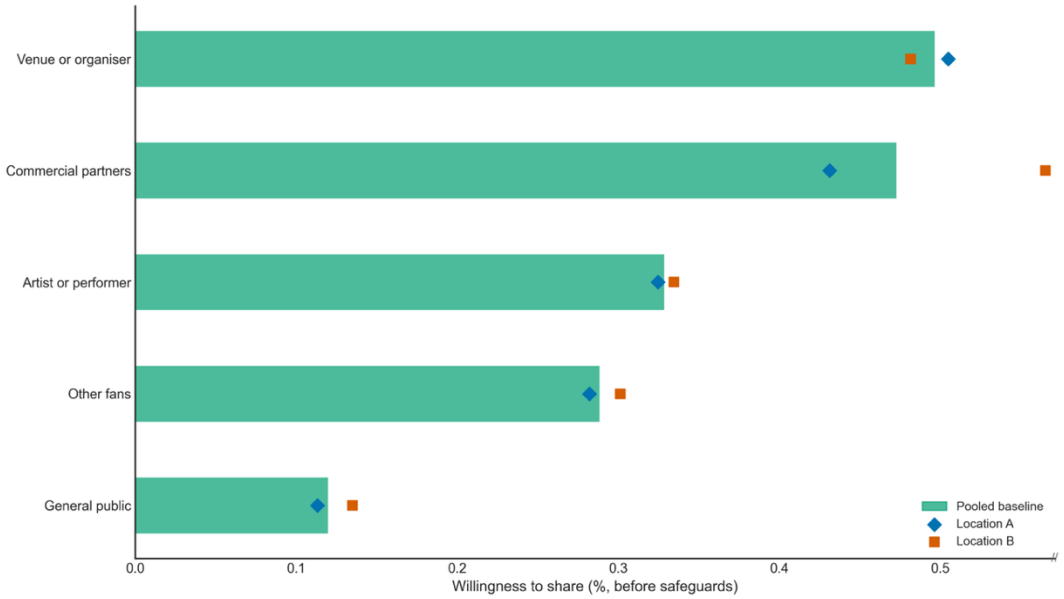

Figure 3. Data-Type Sensitivity Hierarchy.

3.2.3 Near-Consensus Refusal Zones

> *"Location data sharing is usually a very controversial thing, at least for me, because I don't know if it might have an option to just show your location to the people that you prefer rather than everybody else in the fan community or at the event space, because sometimes you don't want that." FG01-P04*

89 of 275 scenarios (32%) show near-unanimous refusal before any safeguards – fewer than 1 in 5 participants willing to consent in each. These 'red-zone' cells are not anomalies. They represent structurally patterned resistance that governance must contend with, not isolated edge cases.

The red zone concentrates around two structural patterns: Diffuse publics (General public accounts for 44 red-zone cells, Other fans for 26, Artist/performer for 19; institutional stakeholders – Venue and Commercial partners – carry far fewer); and sensor-rich data at high-traffic moments, where the heaviest concentration sits on `Venue check-in x Device data` (wristband ID, sensor readings).

Figure 4 maps baseline willingness as a heatmap across all 275 scenarios, with cells coloured by willingness level and the 89 red-zone cells (≤20%) concentrated in the general public row and account/device data columns. Darker cells indicate lower baseline willingness. The general public row and account/device data columns concentrate the deepest refusal – the 89 red-zone cells sit overwhelmingly in these regions.

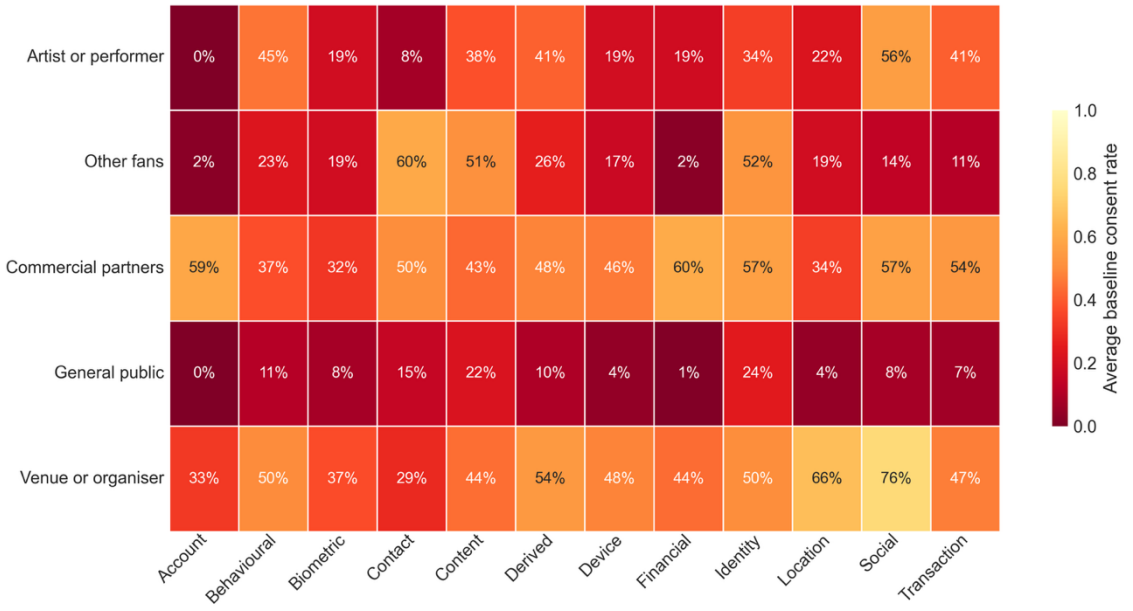

Figure 4. Near-Consensus Refusal Zones.

3.2.4 Within-Sample Heterogeneity

There is no average participant. Participants varied considerably in their overall openness to data sharing. The middle 50% of the sample spans a 14-point range, meaning the typical participant is far from representative of the full spread – see Appendix A.4.2 for details of our Breadth of Sharing (BSI) metric. This heterogeneity is real rather than an artefact of pooling across locations: The two groups show moderate between-location variation, so the pooled picture should be read with a location qualifier. Figure 5 shows the distribution of BSI scores across the full sample.

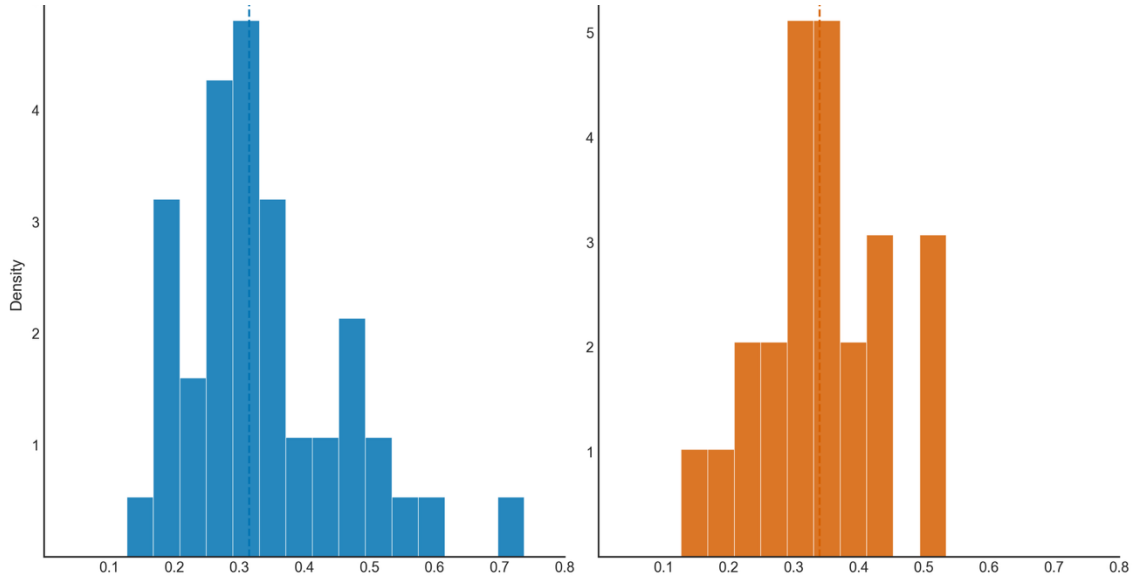

Figure 5. BSI Distribution (Location A / Location B).

3.3 What Governance Unlocks – and Where It Fails

> *'I think I felt more comfortable with it once the consent things were put in place, so I wouldn't mind my data being collected as long as I just have control over it and I understand as well what it means, not a bunch of jargon and stuff that's trying to trick you or whatever, sort of clear communication and easily being able to control that.' (FG08-02)*
>
> *'I am very hesitant with sharing data. If there's a way to opt out of something, I'd rather opt out of it than have my data used.' (FG03-P03)*

Governance improves willingness in only 34% of scenarios; two-thirds remain resistant regardless of which mechanism is applied. No scenario showed net withdrawal of consent after governance. Governance does not reduce willingness in aggregate – it simply fails to move it in two-thirds of the matrix.

After governance mechanisms were introduced, the consent matrix was re-administered. Across all 275 scenarios (pooled, n = 70):

- High lift (≥20 percentage points): 16 scenarios (5.8%) – a strong governance response.
- Moderate lift (10 –19 percentage points): 77 scenarios (28.0%) – a meaningful but partial response.
- Resistant (<10 percentage points): 182 scenarios (66.2%) – limited or no response to governance.

Maximum lift observed was +33.3 percentage points, in the `Avatar creation x Biometric data x Venue` or organiser combination. Median lift across all scenarios was +7.5 percentage points. Figure 6 shows the distribution of all 275 scenarios across the three unlockability tiers, making the scale of the resistant mass visible at a glance.

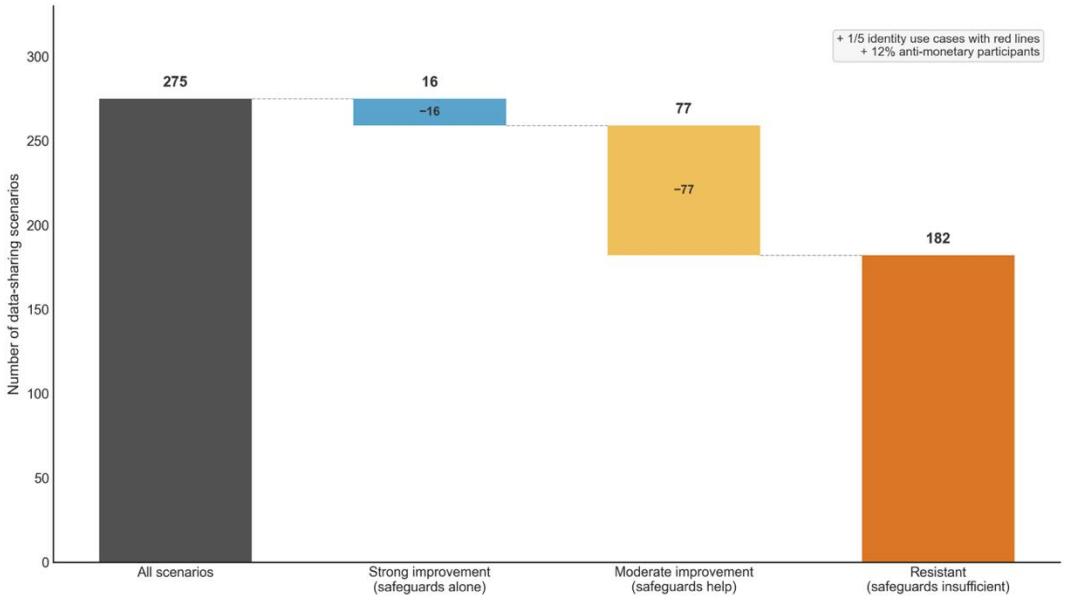

Figure 6. Governance Effect: Unlockability across all Scenarios.

Figure 7 shows governance lift for each of the 275 scenarios individually, sorted from highest to lowest, with the tier thresholds marked and the top and bottom ten cells labelled.

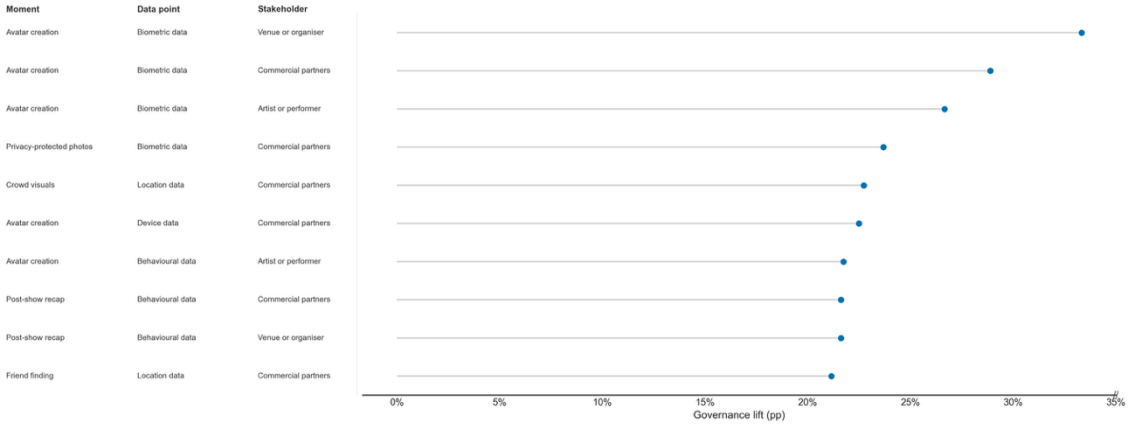

Figure 7. Governance Lift: Cell-Level.

3.3.1 The Headline Tension

The experiences participants value most are the ones governance does least for. Both Tier-1 moments are classified as resistant to governance: Accessibility features shows mean lift of +7.3 percentage points against baseline friction of 67.9%; Venue check-in shows +6.6 percentage points against 68.0% friction. Governance renders 0 of the 2 high-priority experiences conditionally feasible on its own.

The tension is sharpest for Venue check-in. This is the most widely valued experience in the study – yet `Venue check-in x Device data` is simultaneously the single heaviest-refusal combination in the entire consent matrix. The resistance is not to frictionless entry but to the data infrastructure that enables it – wristband device data, real-time sensor readings etc. Participants want the outcome, but they resist the mechanism.

Table 1 shows all 14 moments sorted by governance lift, making this pattern immediately visible: The two Tier-1 moments (highlighted) sit near the bottom of the lift ranking, while the four governance-enabled moments are all lower-value experiences. Baseline friction (RFS) shows how much resistance governance has to overcome. The combination of high RFS and low lift – visible in the two Tier-1 rows – is the headline tension: The experiences people value most are the hardest to unlock.

Experiential moment	Value tier	% Selected	Baseline friction (RFS)	Mean governance lift	Unlock tier	Configuration
Avatar creation	Contested	13.0%	71.0%	+15.5%	Moderate	Governance-enabled
Post-show recap	Tier 2	33.3%	69.3%	+12.2%	Moderate	Governance-enabled
Privacy-protected photos	Tier 2	42.0%	66.3%	+10.9%	Moderate	Governance-enabled
Crowd visuals	Contested	15.9%	64.3%	+10.6%	Moderate	Governance-enabled
Remote hologram	Contested	18.8%	61.9%	+9.0%	Resistant	Governance-resistant
Artist & venue analytics	Tier 2	58.0%	61.4%	+9.0%	Resistant	Governance-resistant
Digital marketplace	Contested / Niche	18.8%	62.3%	+7.3%	Resistant	Governance-resistant

Accessibility features	Tier 1	66.7%	67.9%	+7.3%	Resistant	Governance-resistant
Social messaging	Tier 2	36.2%	69.3%	+7.2%	Resistant	Governance-resistant
Friend finding	Tier 2	49.3%	67.9%	+7.2%	Resistant	Governance-resistant
Venue check-in	Tier 1	58.0%	68.0%	+6.6%	Resistant	Governance-resistant
Audience voting	Contested	30.4%	55.8%	+6.5%	Resistant	Governance-resistant
Content licensing	Contested	21.7%	66.5%	+5.4%	Resistant	Governance-resistant
Wearable setup	Tier 2	37.7%	68.7%	+4.8%	Resistant	Governance-resistant

Table 1. Governance Feasibility by Experiential Moment

Red rows = Tier-1 priority moments. Green rows = governance-enabled moments. RFS = Relative Friction Score (proportion of the consent matrix set to "no" at baseline). Governance lift is the mean change in willingness across all cells within the moment after governance mechanisms were introduced.

3.3.2 Controls that Enable Acceptance

> *"I did the anonymisation one, just because if it's anonymous it is still my data but can't be traced back to me specifically . . . if it's anonymous, then it's fine for some things." (FG09-P02)*
>
> *'I feel like that's just like a safety barrier, and like you can take away your consent at any given moment, which is kind of cool.' (FG04-P05)*
>
> *'Something like the limited time consent, I think would be good if it reset after every concert.' (FG16-P04)*

Three mechanisms account for 56.4% of all consent changes: Anonymisation safeguards (21.4%), time-limited/revocable consent (19.0%), and verified technical controls (16.0%). The remaining seven mechanisms each contribute less than 10%.

These three mechanisms share a structural property: Each gives participants confidence that what is promised is demonstrably enforced – not just stated policy, but provable compliance. A plausible interpretation of why these three lead the rankings is that verified technical controls means cryptographic proofs, independent audits, or certification labels confirming that anonymisation, retention limits, and stated policies have actually been enforced, not merely promised. The Mechanism Salience Score (MSS) used here measures the share of No to Yes transition credits attributed to each mechanism, using a credit-splitting rule where multiple mechanisms were cited.

Figure 8 shows the MSS for each of the 10 governance mechanisms pooled, with location comparison overlaid, confirming the three-mechanism lead and the gap between the top tier and the remainder.

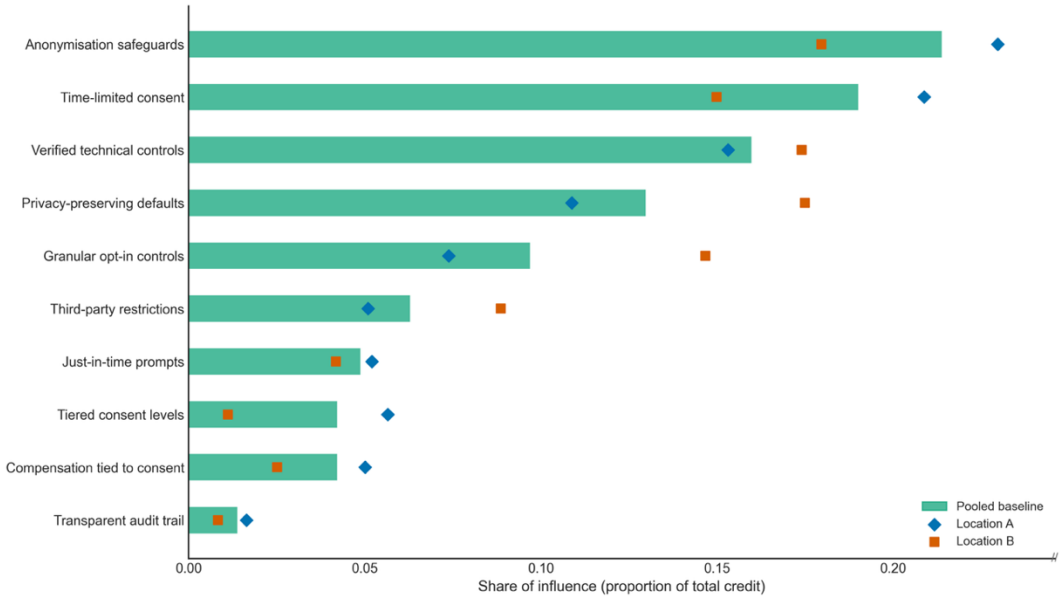

Figure 8. Mechanism Salience Score (MSS).

3.4 Fairness and Compensation

> *'I put revenue shares number one because like honestly it saves like whatever they use my data for goes viral, they're accumulating way more cash than direct cash payment because they don't just give me a fixed amount of money but there there's like my revenue is …and it has no limits.' (FG04-P01)*
>
> *'I chose direct cash payments. I think the main thing for me was like it's a direct kind of financial outcome like it can be measured in terms of how much your clip is used or what you've contributed.' (FG13-P04)*

Monetary compensation dominates strongly: Revenue share and direct cash payments are both selected by 87.7% of participants, fairly well clear of all other options. These two mechanisms are co-selected by 78% of participants – fairness expectations are composite, not singular. Participants want both one-off and ongoing monetary recognition, not one or the other. Exclusive access or perks (VIP) is the strongest non-monetary mechanism at 69.2%. Participants ranked compensation mechanisms using the same Top-5 protocol (n = 65, pooled).

Figure 9 shows selection rates for each compensation mechanism sorted from highest to lowest, with location comparison overlaid and the gap between the two top-tier mechanisms and the rest clearly visible.

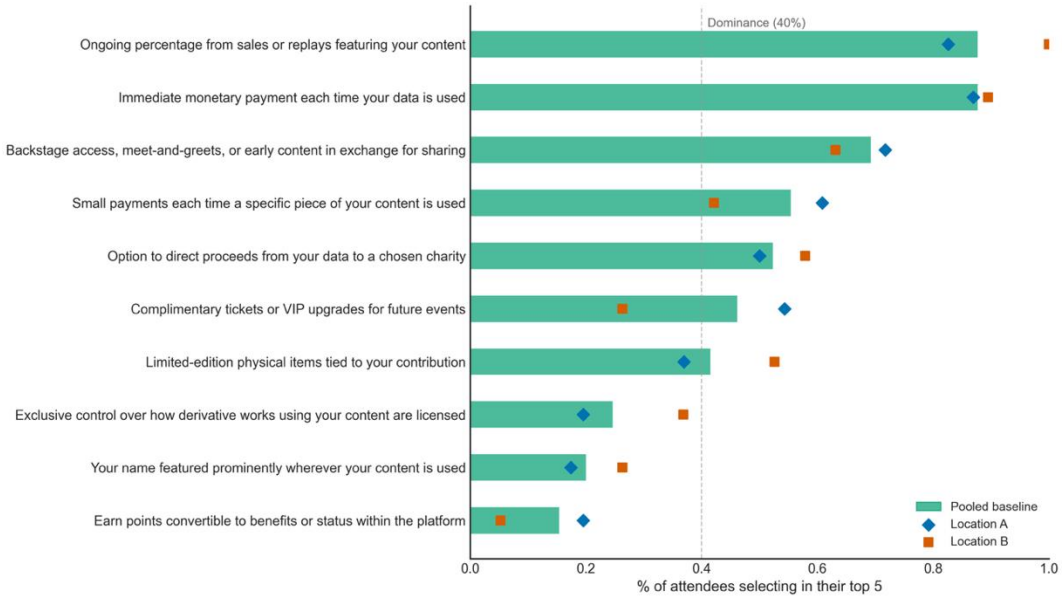

Figure 9. Compensation Mechanisms Ranking.

3.4.1 Compensation as Complement, not Substitute

Compensation does not resolve resistance where governance fails; it complements governance but cannot substitute for it. Cross-referencing governance unlockability with compensation demand across all 14 moments reveals two patterns: Governance-sufficient (4 moments – governance lifts acceptance with no compensation override); and Neither-sufficient (10 moments – governance, compensation, and identity accommodation all insufficient). No moment sits in a 'both needed' quadrant at the defined thresholds.

It is important to note that because compensation was ranked globally rather than per moment, compensation demand is effectively constant across all 14 moments; the cross-referencing is therefore driven by governance tier. Compensation works alongside governance where governance is working; it cannot substitute for governance that is not.

Figure 10 plots all 14 moments on axes of governance tier (x) and compensation demand (y), with the clustering of moments in the top-left quadrant – compensation demanded but governance insufficient – reflecting the dominant pattern in the data.

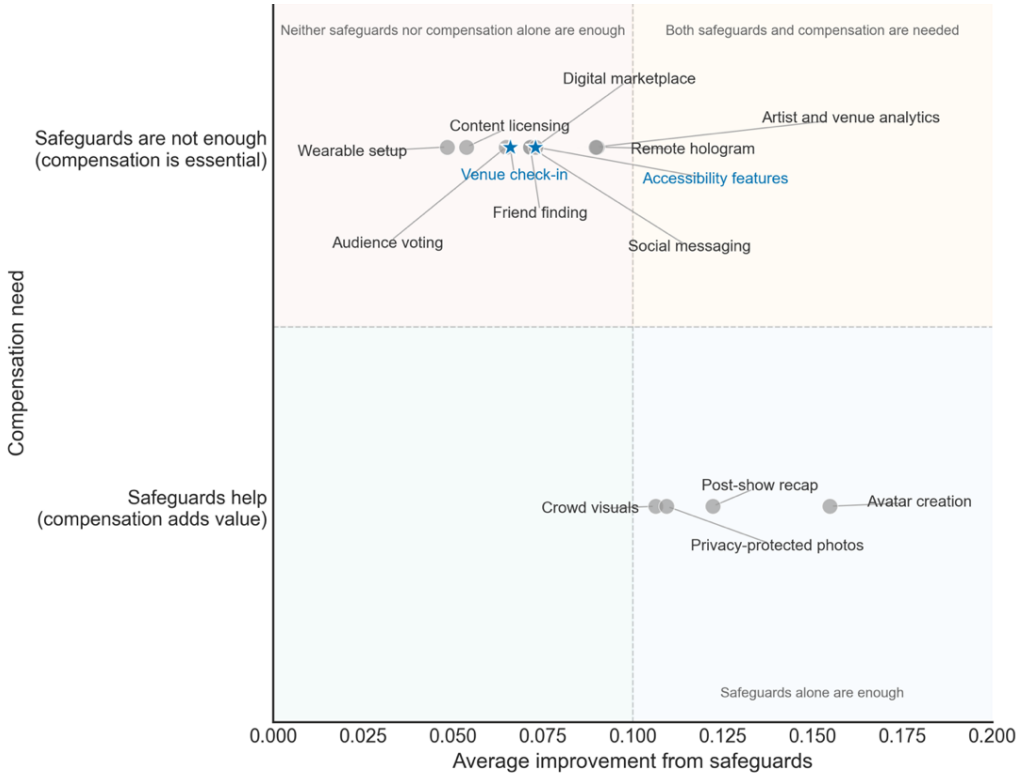

Figure 10. Governance / Compensation.

3.4.2 Rejecting the Market Framing

12.3% of participants – approximately 1 in 8 – selected zero monetary compensation mechanisms, treating data and likeness as outside the scope of market exchange. This is a minority position, but normatively significant as it challenges the assumption that compensation can serve as a universal governance complement. Any consent architecture premised on compensation as a legitimation strategy will have inherent coverage limits and will systematically fail this group. These participants framed compensation as either governance-based (enhanced control, revocation rights) or rejected it in principle.

3.5 Identity Boundaries

The identity and likeness exercise reveals a domain that operates under fundamentally different logic from informational data sharing. Participants indicated their preferred representation level (Real identity, Avatar, or None/Refuse) for each of five identity use cases (see Table B.4 in Appendix B), both before and after governance mechanisms were discussed.

3.5.1 Data Sharing and Identity Representation are Distinct

> *'An avatar, I think, gives you that protection that you don't have to be yourself in the world, but then you can still be a part online. But then again, if you are a creator yourself, I think you would want your real image in all of these.' FG13-P02*

Governance unlocks identity representation more readily than data sharing – a counter-intuitive finding that indicates a need for separate governance architectures for the two domains. Identity willingness is consistently higher than equivalent data-sharing willingness across all five use cases, ranging from 16.9 percentage points higher (Personal highlight reel) to 39.4 percentage points higher (Promotional material). Treating consent as a single dimension misses this distinct logic.

The two dimensions move independently: Governance that works for data does not automatically transfer to identity, and vice versa. The likely explanation is that the five identity use cases are more purposeful and bounded than open-ended data collection, making them easier for participants to assess against their own preferences.

A further explanation emerged from the qualitative data, namely that some participants viewed avatars as commercially separable from their real identity, treating their use as a monetisation opportunity rather than a privacy compromise. This framing may have contributed to the size of the avatar uptick observed in the more commercial use cases. The Avatar Buffer Index should therefore be read as capturing two distinct movements: (1) Avatar as protection against real-likeness exposure, and (2) avatar as an acceptable monetisable proxy.

> *'Yeah, for the avatar merchandise, just because, money.' (FG04-P04)*
>
> *'I chose number 2 because using my avatar doesn't seem that invasive at the end of the day, it's just an avatar, I can create it in any way I want to and there is a tangible profit out of it, 10% net sales, so it's good.' (FG08-P06)*

Table 2 below shows the comparison for each identity use case. 'Data-sharing lift' is the mean governance lift for data-sharing scenarios associated with that use case (8.5 percentage points pooled reference). 'Identity upgrade rate' is the proportion of participants who moved to a higher representation level after governance. The 'Gap' is the difference between identity upgrade rate and data-sharing lift – positive values show identity is more governance-responsive. 'ABI' (Avatar Buffer Index) measures whether avatar selection increased after governance: Positive means avatars absorbed some refusals as a compromise. A red-line flag indicates that ≥30% of participants still choose None post-governance.

Identity use case	Identity upgrade rate	Gap (identity minus data)	ABI	Red line
Promotional material	+39.4%	+30.9 pp	+0.14	–
Avatar in virtual spaces	+28.2%	+19.7 pp	+0.07	⚑ Yes (35.2%)
AI after-movie	+28.2%	+19.7 pp	0.00	–
Social media content	+22.5%	+14.0 pp	0.00	–
Personal highlight reel	+16.9%	+8.4 pp	−0.04	–
Pooled data-sharing reference	–	8.5 pp	–	–

Table 2. Data Sharing vs Identity Consent by Use Case. Red row = Avatar in virtual spaces carries the red-line flag (35.2% persist in None post-governance). Gap = Identity upgrade rate − pooled data-sharing lift (8.5 pp, the median lift across all 275 data-sharing scenarios). ABI = Avatar Buffer Index (%Avatar post − %Avatar pre).

3.5.2 Avatar Mediation as Compromise

Avatar mediation functions as a partial compromise pathway, but with effectiveness varying substantially by use case (ABI range: −0.04 to +0.14) and without reaching the segment that objects to virtual-space presence itself. Of 355 total identity transitions,

the dominant flows show stability: Real to Real (101 cases), None to None (86), Avatar to Avatar (68). However, of participants initially in the 'None' state, 73 committed to an identity representation after governance: 42 moved to Real and 31 to Avatar. The 86 who remained in 'None' indicate an absolute-refusal segment that governance, compensation, and avatar mediation did not reach.

Figure 11 shows identity preference flows across all five use cases pooled, with ribbon width proportional to the number of participants following each pre-to-post transition path.

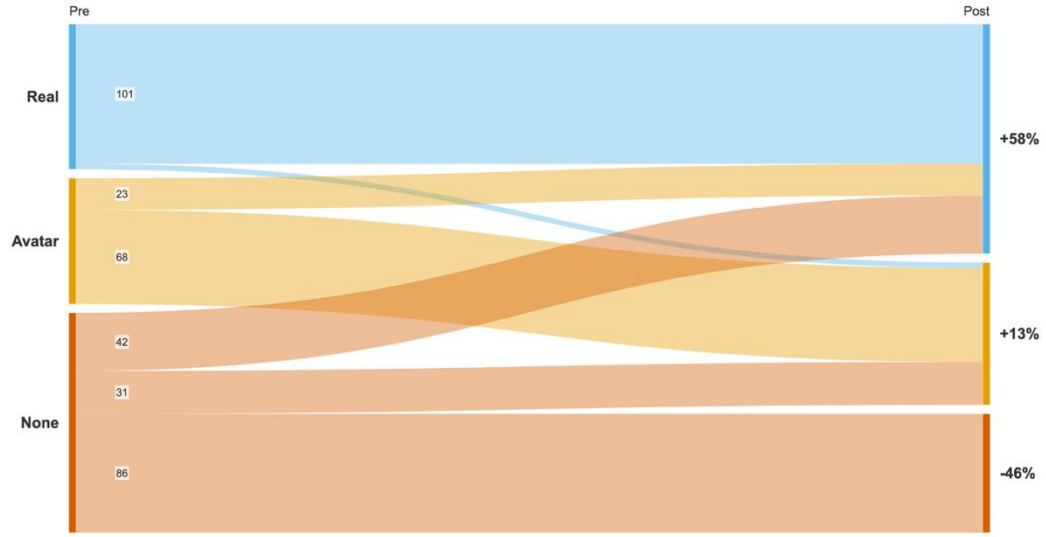

Figure 11. Identity Transition Flows.

Table 3 shows how the proportion of participants choosing Real identity, Avatar, and None (refuse) shifts from before governance (Pre) to after governance (Post) for each use case.

Identity use case	Real (Pre)	Avatar (Pre)	None (Pre)	Real (Post)	Avatar (Post)	None (Post)	Red line
Promotional material	15.5%	18.3%	66.2%	38.0%	32.4%	29.6%	–
Avatar in virtual spaces	8.5%	33.8%	57.7%	23.9%	40.8%	35.2%	⚑
AI after-movie	39.4%	19.7%	40.8%	60.6%	19.7%	19.7%	–
Social media content	62.0%	15.5%	22.5%	76.1%	15.5%	8.5%	–
Personal highlight reel	22.5%	40.8%	36.6%	35.2%	36.6%	28.2%	–

Table 3. Identity Representation Before and After Governance

Values expressed as % of participants (n = 71). Red row = red-line threshold breached (≥ 30% None post-governance). Avatar in virtual spaces is the only use case to cross this threshold.

3.5.3 Persistent Identity Red Lines

> *"I feel like . . . [if an] avatar could walk on the stage . . . it would be a bit distracting like visually, and like you're going to see a band or like a real people, you don't want like little things jumping up everywhere." FG02-P02*

One of five identity use cases carries a hard limit that governance, compensation, and avatar mediation cannot resolve: Avatar Representation in Virtual Spaces, where 35.2% of participants still refuse post-governance. The apparent paradox – avatars partially working yet a red line persisting – resolves when the use case is understood to contain two distinct sub-populations. Those who object to real-likeness exposure found the avatar option workable; those who object to virtual-space digital presence in any form could not be reached by avatar mediation, because the avatar remains an instantiation of the thing they object to. Compromise and red line coexist because they belong to different people.

Identity red lines are categorically distinct from reluctance over data sharing. Reluctance is a matter of degree (more or less willing); identity red lines are categorical (unacceptable regardless of safeguards). Governance is designed for reluctance; it has no purchase on red lines.

Figure 12 shows the change in outright refusal rate (%None) for each identity use case from pre- to post-governance, with Avatar in Virtual Spaces flagged as the single use case crossing the red-line threshold.

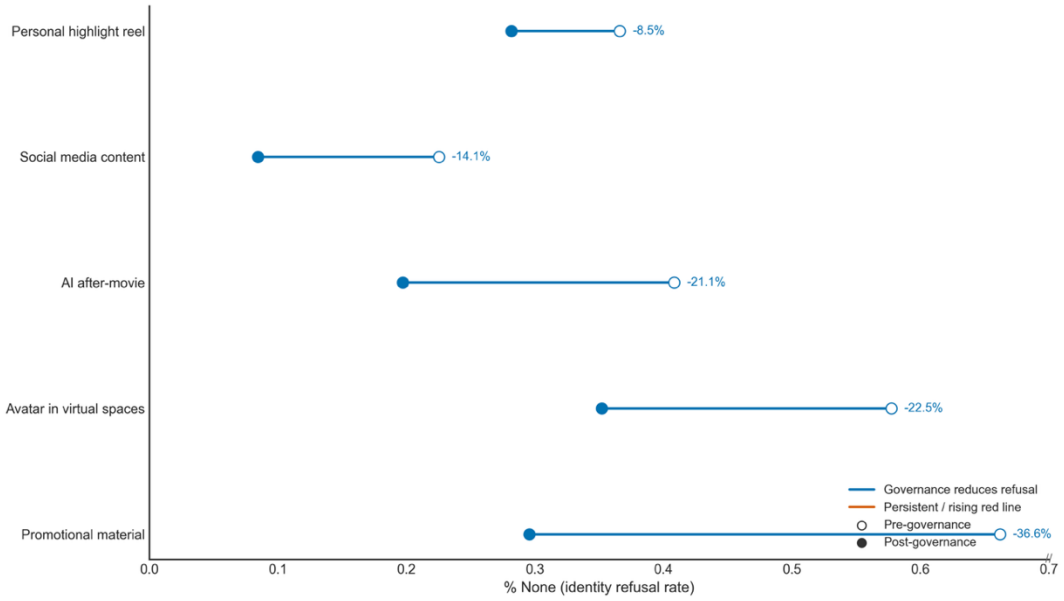

Figure 12. Identity Divergence.

3.6 Configurations and Structural Limits

Integrating across all exercises, we coded each of the 14 moments across four dimensions – value tier, governance unlockability, compensation stance, and identity constraint – to identify recurring configuration patterns.

3.6.1 The Configuration Typology

The 14 moments resolve to two empirical configurations: Governance-enabled (4 moments) and governance-resistant (10 moments), with both Tier-1 value moments sitting in the harder group. The governance-resistant moments score higher on participant value importance – the harder cases are harder partly because the stakes are higher.

The conceptual framework distinguishes four types, each with distinct design implications:

- Governance-sufficient: High governance lift, no dominant compensation demand, identity unconstrained or avatar-buffered. The tractable design cases where procedural safeguards do the job.
- Fairness-required: Governance lifts acceptance meaningfully, but participants simultaneously demand monetary fairness. Governance is necessary but not sufficient; compensation closes the gap.
- Identity-buffered: Governance alone does not resolve resistance, but avatar mediation provides a compromise pathway.
- Governance-resistant: Governance, compensation, and avatar mediation are all insufficient. A structural limit.

The four governance-enabled moments are: Avatar creation, Crowd visuals, Privacy-protected photos, and Post-show recap. The ten governance-resistant moments are: Venue check-in, Accessibility features, Wearable setup, Social messaging, Friend finding, Audience voting, Remote hologram, Digital marketplace, Artist and venue analytics, and Content licensing.

Figure 13 shows all 14 moments classified by configuration type, sorted by mean governance lift, with the two Tier-1 moments visible within the governance-resistant group at the foot of the lift ranking.

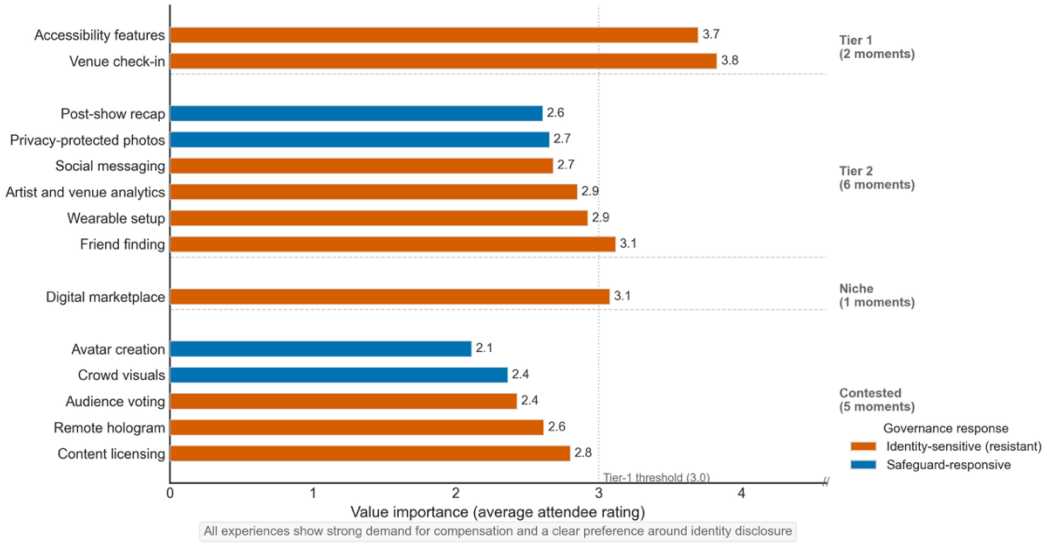

Figure 13. Configuration Typology.

3.6.2 Three Categories of Structural Limit

Three distinct categories of structural limit are identified, none of which represents a failure of implementation. Rather, they indicate domains where technical and procedural solutions may be insufficient as a matter of principle:

- Scenarios where willingness did not respond to governance: 182 of 275 (66.2%) show lift below +10 percentage points regardless of mechanism. (The 275 scenarios refer to specific data-sharing combinations – a `moment x a data type x a stake`holder – not to the 14 moments or 5 identity use cases.) Governance alone fully unlocks only 16 scenarios and partially improves another 77.

- 1 of 5 likeness use cases carries a red line – ≥30% of participants persist in refusal post-governance even after all safeguards are applied. Where this threshold is crossed, the boundary is categorical, not incremental.
- Anti-commodification stances: 12.3% of participants reject market-based framing entirely. Any compensation-based legitimation strategy will have inherent coverage limits.

Responsible AI design must recognise these as hard limits, not optimisation targets.

3.7 Location Comparison and Heterogeneity

All four research areas show location divergence between Location A and Location B; no area converges cleanly. The four areas are experience value priorities, change in willingness after safeguards, compensation preferences, and likeness/image preferences. At the consent-matrix level, 71 of 275 combinations (25.8%) diverge between locations.

Both Tier-1 moments are shared across locations, and the stakeholder trust gradient rank order agrees – these foundational patterns are stable. However, 6 of 14 moments are ranked differently between locations, the dominant compensation type differs, and identity dynamics show the largest between-location gap.

Clustering analysis across three dimensions – sharing orientation, value priorities, and post-governance identity refusal – confirms that pooling is justified for value priorities, qualified for sharing orientation, and weakest for identity dynamics (see Appendix A.4.2 for ICC definition). An implication here is that governance frameworks should not assume universal acceptance across different populations.

Figure 14 summarises location stability across the four research areas, showing the proportion of key metrics that converge or diverge between England and Scotland.

Domain	Metric	Diverging Cells	% Diverging
Experience value priorities	PctSelected	6/14	43%
Change in willingness after safeguards	Lift_pp	120/275	44%
Compensation preferences	PctSelected	6/10	60%
Likeness and image preferences	PctNone_Post	4/5	80%

Figure 14. Summary of Divergence across Locations.

4. Implications for Responsible AI Design

4.1 Five Architectural Requirements

Given the findings here, we suggest structural properties any Responsible AI consent system in live event contexts must have – not recommendations about features, but requirements about how the system must be built. Against each point, we suggest an underlying design principle for practitioners:

- Separate consent tracks for data and identity. Because the two dimensions follow different logic and respond differently to the same mechanisms, a unified consent interface will systematically mishandle one or both. For example, a single consent toggle cannot accommodate the categorical distinction participants draw between informational data and personal likeness. *Design principle: Implement two independent consent layers with separate data models, escalation paths, revocation mechanisms, and governance logic. Neither layer should inherit defaults or permissions from the other.*
- Design for enforceability, not just transparency. The three highest-performing governance mechanisms – anonymisation safeguards, time-limited/revocable consent, and verified technical controls – share a structural property worth highlighting as one plausible explanation for their high MSS rankings; each gives participants a defined, bounded, and enforceable basis for their consent. Anonymisation transforms data in ways that cannot be undone; time limits and revocability define when consent ends and what happens to the data when it does; technical certifications provide external proof that stated policies have been followed. What moves willingness is not being told what will happen but being given a structural guarantee that it will. *Design principle: Governance mechanisms should be implemented at the infrastructure level, not the policy level. Anonymisation should be enforced technically, not declared in a privacy notice; consent expiry should be system-enforced, not user-managed; compliance claims should be externally verifiable, not self-certified.*
- Do not rely on compensation to close gaps governance cannot. The data establishes two distinct limits on compensation's reach. First, the data shows no moment for which compensation resolves the residual resistance left by governance. Second, 12.3% of participants reject market-based framing entirely, regardless of how compensation is structured. Compensation is a legitimate fairness expectation for the majority; it is not a governance substitute for anyone, and it will not reach a principled minority regardless of price or mechanism design. *Design principle: Treat compensation as a parallel fairness layer, not a consent mechanism. It should be presented after consent is established, not as a condition of it. Non-monetary consent pathways must be available as a first-class option, not a fallback.*
- Architect around hard limits, not just unlocks. Two-thirds of the consent matrix is structurally resistant. The responsible design question is not 'how do we overcome resistance?' but 'how do we operate ethically within it?' Treating irreducible limits as a conversion problem is itself an architectural failure. *Design principle: The system should be designed to degrade gracefully at consent boundaries. Where a scenario falls in the resistant or red-line zone, the system should default to non-collection or anonymised-only operation rather than seeking consent optimisation. Hard limits should be encoded as system constraints, not as targets for re-engagement.*
- Make heterogeneity a first-order design constraint. Sharing orientations vary too widely for population-level defaults to serve all participants adequately. Personalisation, tiering or meaningful opt-out pathways are structural

necessities, not UX refinements. *Design principle: The consent architecture should support a meaningful range of sharing orientations at the point of onboarding, not converge users toward a single default. Preference profiles should be persistent, portable, and revisable, with the most privacy-protective option as the system default.*

4.2 Implications for P3R WP3 Identity Research

The identity and likeness findings have specific implications for P3R identity research concerned with provenance, attribution and rights management in AI-mediated contexts. The analysis that follows is directly relevant to extending decentralised identity provenance frameworks – such as the ORA (Ownership-Rights-Attribution) model developed under the DECaDE programme – from media content into the identity and likeness domain.

- Provenance tracking requires dual-domain architecture. Data-sharing consent and identity consent follow different logics and respond differently to governance. Provenance systems that conflate the two will be structurally inadequate: A wristband-generated behavioural data trail and an AI-derived likeness use are not the same type of consent transaction, do not share the same governance leverage points and should not be governed by the same mechanism. This has direct implications for applying ORA-style tokenised licensing to identity: The ownership-rights-attribution chain for identity/likeness content should operate on a separate track from data provenance, with distinct consent primitives – not as an extension of an existing data provenance layer.
- Avatar-derived content presents a new attribution challenge for provenance infrastructure. The ABI data shows that avatar representation absorbed refusals in some use cases that neither governance nor compensation could resolve –promotional material (+0.14) and virtual-space contexts (+0.07). This suggests that a functioning identity-provenance system will need to handle avatar-derived content as a distinct category: Proxy-identity content with its own ownership, rights, and attribution chain, separable from real-likeness content. This is a research proposition that the identity extension of the DECaDE programme is positioned to test, building on the TrustMark and metadata-embedding capabilities already developed for media content.
- The two-subpopulation structure in Avatar in Virtual Spaces is an inference that warrants direct investigation. The transition data shows that 16 of the 41 participants initially refusing this use case moved to an identity position after governance, while 25 remained in outright refusal. The most analytically coherent interpretation is that these represent two distinct sub-populations with different underlying objections –
one negotiable (objection to real-likeness exposure), one categorical (objection to virtual-space digital presence itself). However, the current data does not include attribution of refusal to specific objection types – this is an inference, not a measurement. For identity-provenance research, this distinction is consequential: The ORA framework can potentially serve the first group through proxy-identity mechanisms; the second group is outside the reach of any licensing model, regardless of granularity. Future research should test this interpretation directly through attribution-linked refusal data.
- The ORA micropayment model faces a principled ceiling of approximately 12% in this population. The 12.3% anti-commodification rate is broadly stable across locations (13.0% at Location A, 10.5% at Location B) and is independent of compensation structure or granularity: These participants selected zero monetary mechanisms of any kind. For identity provenance research extending

ORA's micropayment mechanism into the identity domain, this finding sets a coverage ceiling: roughly 1 in 8 participants in this sample will not engage with any market-based legitimation strategy, including per-clip licensing, per-replay royalties, or NFT-based rights assertion. This does not invalidate the ORA model for the majority; it establishes that a functioning identity-rights system built on that model requires non-monetary consent pathways alongside it, not as an afterthought.

- Governance lift for identity is use-case-specific, not category-general. ABI varies from −0.04 (Personal highlight reel, where governance moves people past the avatar option entirely) to +0.14 (Promotional material). Identity governance interventions cannot be designed at the category level – 'identity use' is not one thing. For the ORA framework applied to identity, this means that the governance and licensing architecture appropriate for promotional material use of a likeness will differ substantially from that appropriate for virtual-space avatar deployment of the same person's identity. Research programmes should calibrate governance mechanisms per use case and test effectiveness at that level of granularity.
- The red-line threshold provides a replicable benchmark for identifying use cases outside the scope of tokenised licensing. The 30% post-governance refusal threshold used here identifies use cases where categorical resistance is present – where the objection is not to the terms of a licence but to the existence of the use case itself. For ORA extended into identity, this threshold is practically useful: use cases that cross it should be excluded from the scope of a tokenised licensing framework, not designed around, because no licensing mechanism – however granular, transparent or fairly-priced – can reach a categorical objection. Use cases that sit near but below the threshold (e.g., Promotional material at 29.6% None post-governance) represent boundary cases warranting careful monitoring.

Appendix A: Methodology

A.1 Research Design

WP2 employs a design fiction methodology. A video depicting a future AI-enabled live music event was created as a provocation, showing 14 distinct experiential moments involving various forms of data collection, AI processing, and personalisation. Participants were asked to respond to this scenario through a structured sequence of exercises designed to elicit situated judgments about consent, governance, compensation, and identity.

A.2 Participants

Seventy-three active participants were recruited across two UK university locations: Royal Holloway, University of London (Location A, England, n = 48) and Abertay University (Location B, Scotland, n = 25). Participants were self-selecting students with an interest in live music and technology. Demographic composition: 51% identified as male, 41% as female, 8% as non-binary; 88% aged 18–24, 11% aged 25–34. Seventy-three participants attended; two were excluded after data review (one flagged inactive, one with incomplete data mapping), leaving an active sample of 71. (Active sample sizes vary by exercise due to missing data; full per-exercise counts are given in Appendix B.5.)

A.3 Exercise Structure

All focus groups followed an identical protocol comprising six sequential exercises:

- Value perception sprint (15 min): Participants individually selected and ranked their Top-5 experiential moments from the 14 depicted in the Design Fiction, providing short written rationales. A group debrief followed.
- Data portrait mapping (20 min): Each participant completed a baseline consent matrix (14 moments × up to 11 data types × 5 stakeholders = 275 scenarios), indicating willingness to share before any governance mechanisms were introduced.
- Consent shift challenge (20 min): After presentation and discussion of 10 governance mechanisms, participants re-completed the consent matrix. Attribution markings indicated which mechanisms enabled any changes.
- Compensation ranking (15 min): Using the same Top-5 protocol, participants ranked 10 compensation mechanisms, providing rationales for their selections.
- Likeness and image rights (15 min): For each of 5 identity/likeness use cases, participants indicated their preferred representation level (Real, Avatar, or None) before and after governance framing.
- Reflective discussion (15 min): Open-floor group discussion on concerns, red flags, and views on Responsible AI. A final replay-and-reflect exercise tested whether opinions had shifted during the session.

A.4 Analytic Approach

Our analysis was governed by a detailed Analysis Playbook establishing procedural consistency across all focus groups, ensuring outputs are traceable, comparable, and methodologically defensible.

A.4.1 Data Architecture and Quality Assurance

All data were entered into a standardised workbook using locked master identifier lists. All derived fields were system-generated. Missing data were never treated as negative responses: A blank cell represents unknown willingness, not unwillingness. Imputation was not permitted. Structured validation was conducted for each exercise prior to analysis, and all deviations were recorded in a dedicated log.

A.4.2 Key Metrics and Thresholds

Key metrics include:

- Percentage willingness (%Yes_pre, %Yes_post) – the share of participants willing to consent to a given data-sharing combination. Calculated at the cell level (moment × data type × stakeholder) as the number of yes responses divided by the number of valid responses for that cell, expressed as a percentage.
- Governance lift (Lift_pp) – the change in willingness attributable to the introduction of governance safeguards. Calculated as %Yes_post minus %Yes_pre, and classified into unlockability tiers: High (≥20 pp), Moderate (10-19 pp), and Resistant (<10 pp).
- Mechanism Salience Score (MSS) – the relative importance of each governance mechanism in driving consent. Calculated as the proportion of all No to Yes transition credits attributed to each mechanism, using a credit-splitting rule where multiple mechanisms were cited together.
- Breadth of Sharing Index (BSI) – each participant's overall openness to data sharing across the full exercise. Calculated as the proportion of all valid cells in which that participant consented (0 = shares nothing; 1 = shares everything).
- Avatar Buffer Index (ABI) – the extent to which avatar representation functions as a compromise pathway in identity contexts. Calculated as %Avatar_Post minus %Avatar_Pre across pre- and post-governance conditions.
- Identity red-line flags – use cases representing categorical refusal that governance does not resolve. Flagged where ≥30% of participants persist in refusal post-governance.
- Intraclass Correlation Coefficient (ICC) – the degree to which between-location differences, rather than between-individual differences, account for variation in a metric. Calculated across three measures (BSI, weighted value score, and post-governance identity refusal) to assess whether pooling across locations is justified.

A.5 Interpretive Constraints

The focus groups use design fiction and sequential exercises that intentionally introduce concepts (governance mechanisms, compensation models). Changes in willingness between baseline and governance-framed conditions represent conditional attitudinal shifts under deliberative discussion, not experimentally isolated treatment effects.

Results should not be interpreted as population estimates, behavioural predictions, market demand indicators, or causal effects of governance mechanisms. The sample (n = 70 for baseline analyses, n = 71 for identity analyses) is drawn from self-selecting students at two UK universities with an interest in live music and technology. Focus groups are logistical units, not theoretically distinct cases – location-level comparison is the appropriate unit for examining heterogeneity.

Appendix B: Reference Tables

The following tables list all items used in the focus group exercises.

Table B.1: Experiential Moments (A–N)

Data types referenced across moments include: Identifying data, contact data, demographic data, account and login data, financial data, online/device data, location data, biometric data, behavioural data, communications and content data, transactional data, and derived/inferred data.

Moment	Plain English Description	Data Type Examples
A – Wearable setup (00.16)	Setting up your ConsentWave wristband and connecting it to your account	Identifying data; Account and login data; Online/device data; Financial data
B – Social messaging (00.25)	Sending videos, photos, or texts to friends through the app	Communications and content data; Identifying data; Online/device data
C – Avatar creation (00.45)	Designing your digital character and having it copy your real movements	Biometric data; Behavioural data; Identifying data; Online/device data
D – Digital marketplace (00.59, 01.21)	Buying, selling, or trading digital items; earning credits from your creations	Financial data; Transactional data; Identifying data; Communications and content data
E – Venue auto check-in (01.34)	Your wristband automatically registers your arrival at the concert venue	Identifying data; Location data; Online/device data; Derived data
F – Crowd visuals (01.49)	Your movements contribute to the light show and visuals on stage	Behavioural data; Location data; Biometric data; Derived data
G – Accessibility features (02.23)	Turning on captions, adjusting lights, or enabling vibration feedback	Demographic data; Behavioural data; Online/device data; Derived data
H – Friend finding (01.55)	Getting notified when online friends are nearby at the event	Location data; Contact data; Identifying data; Derived data
I – Privacy-protected photos (02.35)	Taking selfies where other people's faces are automatically replaced with their avatars	Communications and content data; Biometric data; Identifying data; Derived data
J – Audience voting (02.45)	Voting with your wristband to control the lighting, tempo, or visuals of the show	Behavioural data; Online/device data; Derived data
K – Remote hologram (02.59)	Attending virtually and appearing as a life-size hologram at the venue	Communications and content data; Behavioural data; Identifying data; Online/device data
L – Post-show recap (03.21)	Receiving a custom video of the concert based on your experience and viewing choices	Behavioural data; Communications and content data; Derived data; Identifying data
M – Analytics for artists and venues (03.38)	Anonymised crowd data helping performers and venues improve future shows	Derived data; Behavioural data; Transactional data; Demographic data

N – Content licensing and reward (04.00)	Getting rewarded when your photo or content is used commercially	Communications and content data; Identifying data; Financial data; Transactional data

Table B.2: Consent/Governance Mechanisms (1–10)

#	Mechanism	Description	Example Application
1	Granular opt-in controls	Separate toggles for distinct uses (live visuals, promos, merch, analytics, third-party sharing)	A wristband menu lets you separately enable "live visuals", "promo photos", and "third-party analytics"
2	Time-limited consent and easy revocation	Consents expire and can be revoked with clear downstream effects	Allow visuals for 'this event only' with a revoke button that removes future uses from the archive
3	Contextual just-in-time prompts	Brief, actionable prompts at moment of use showing immediate consequences	A one-tap prompt asks permission the moment a photo with your avatar is about to be shared
4	Anonymisation and aggregation safeguards	Enforced de-identification, minimum group sizes, and technical guarantees (e.g., k-anonymity)	Movement data is 'fuzzed' and combined with at least 50 other attendees before use for analytics
5	Transparent audit trail and logging	User-accessible history showing when, how, and by whom their data/likeness was used	A dashboard shows timestamps and recipients each time your likeness or data was used
6	Clear compensation tied to consent	Explicit, proportional rewards shown at consent time	Grant permission for a merchandise drop and immediately receive compensation
7	Default privacy-preserving settings	Privacy-friendly defaults with explicit upgrade paths	New users default to 'no public sharing' with optional switches to enable broader sharing for rewards
8	Tiered consent levels with mapped benefits	Defined consent tiers mapping to concrete benefits and data scopes	Basic = anonymised analytics + small credits; Enhanced = promo use + larger credits; Full = promo + merch + revenue share
9	Third-party use restrictions and contractual limits	Explicit bans or constraints on resale, further sharing, or training commercial models without new consent	A clause that forbids resale of your avatar likeness to external advertisers without a fresh consent prompt
10	Verifiable technical controls and certifications	Cryptographic proofs, independent audits, or labels certifying that anonymisation, retention limits, and stated policies have been enforced	An independent audit badge on the event page confirming that anonymisation and retention policies were enforced

Table B.3: Compensation Mechanisms (1–10)

#	Mechanism	Description	Example
1	Direct cash payments	Immediate monetary payment per use/sale/licensing	ConsentWave pays you £X each time your holographic clip is licensed for a replay
2	Revenue share/royalties	Ongoing percentage from sales/replays/merch where the user's content is used	You receive 10% of net sales whenever merchandise featuring your avatar is sold
3	Platform credits/discounts	Redeemable credits for tickets, merch, or premium features	Grant promo use of your avatar and get 100 platform credits redeemable for tickets or merch

4	Exclusive access/perks	Tickets, backstage access, meet-and-greets, early content access	Allow backstage hologram placement and receive a free VIP ticket plus a meet-and-greet pass
5	Micro-payments for micro-uses	Small payments for specific short uses (e.g., single replay clip)	Each time a 10-second clip of your dance is downloaded, you get £0.XX instantly
6	Exposure and attribution	Prominent crediting, featured placement, or promotion to fan communities	Your username and profile are featured on the livestream credits and promoted to the fan community
7	Physical merch or limited editions	Real-world items tied to the user's contribution	Consent to use your move and get a limited-run T-shirt with your avatar printed and shipped to you
8	Points and gamified rewards	Points convertible to benefits or status within the platform	Earn 250 XP points toward status tiers each time your motion data contributes to a live visual
9	Charitable donations	Option to direct proceeds to a chosen charity	Opt to donate proceeds from uses of your likeness to a selected charity and receive a donation receipt
10	Enhanced control/rights	Options like exclusive control over derivative uses, licensing windows, or veto rights in exchange for reduced monetary reward	Choose a licensing window (30 days) for commercial uses of your avatar in exchange for a lower immediate fee

Table B.4: Identity/Likeness Use Cases (1–5)

#	Image/Likeness Use	Example
1	POV selfie/fan photo posted to feeds	Attendee photos uploaded to social feeds can be reposted, licensed for livestream/album art, sold as NFTs or prints, and/or used in promotional campaigns
2	Holographic/AR recordings of attendees/avatars in mixed-reality streams	Live MR captures and recorded streams (including avatar projections) can be sold as replays, clips, or licensed for highlights and promotional videos
3	Avatar/gesture-derived purchasable content (dance moves, skins, patches)	Motion-captured performances and custom avatar appearances can be packaged and sold as downloadable moves, skins, or collectables in marketplaces
4	Merchandise using attendee images or avatar likenesses	Digital images or avatar renders printed on physical merch (T-shirts, capsules) or offered as limited-edition items that generate resale/royalty revenue
5	Curated/featured content selected for official releases or promos	Photos or clips (e.g., a selfie chosen for an album cover or highlight reel) become part of official media, monetised via streams, physical releases, or sponsorship placements

Table B.5: Participants Count by Exercise

Exercise	Total n	Location A	Location B
Value perception sprint	70	47	23
Data mapping — baseline (pre)	71	46	25
Consent shift — re-mapping (post)	70	46	24
Compensation ranking	66	46	20
Image rights — pre	71	47	24
Image rights — post	71	47	24
Active after exclusions (reference)	*71*	*47*	*24*

Attended (reference)	**73**	**48**	**25**

Note: Variation between exercises reflects missing data, not different recruitment; a participant present in one exercise sheet may be absent from another if that exercise's data is incomplete or invalid.