

Original Paper

From Regulatory Vacuum to Socially Shaped Governance: A Responsibility-Gap Model of Facial Recognition

Qiuting Zhang¹, Ziyu Jia¹ & Yuanyuan Li^{1*}

¹ School of Economics and Management, Shandong University of Aeronautics, Binzhou, 256600, China

* Corresponding authors: Yuanyuan Li, E-mail: yyuanli1991@163.com

Received: October 5, 2026

Accepted: October 12, 2026

Online Published: October 20, 2026

doi:10.22158/elp.v9n2p126

URL: <http://dx.doi.org/10.22158/elp.v9n2p126>

Abstract

Facial recognition has transformed unchanging physical characteristics into instruments of economic and governmental power, but regulation has not kept pace with deployment; thus, enterprises are resorting to superficial ethical behavior. We believe that strategic underinvestment is an inevitable consequence of structural regulatory lag, not a moral failure by individuals, and dispersed social actors can help to address this problem. A delayed stock-flow system is employed to construct an ethical signalling game and a polycentric detection framework with correlated nodes for the construction of a formal responsibility-gap model. We apply theory-guided process tracing to the first facial recognition lawsuit in China (2019-2021) and, in conjunction with a five-regime fuzzy-set comparison, use a reproducible open dataset. We have built a socially constructed path of governance, and litigation, media dissemination, judicial decisions and legislative codification are all non-Western regulatory systems. The first is to extend the scope of ethics-washing to include the omission of under-investment, and the second is to differentiate between an institutional responsibility gap and a traditional epistemic gap.

Keywords

ethics washing, omission-based ethics washing, biometric surveillance, responsibility gap, social shaping of technology

1. Introduction

Facial recognition is also a type of technology. It is a part of the body that cannot be altered after display, and it also has commercial and state power; it serves as an entry token for amusement parks,

functions as a surveillance device in public places, and is also a commodity in algorithmic supply chains (Chen, 2026). In doing so, it will change the basic power structure of people and organisations. The biometric body is no longer under one's control; it has been acquired, saved and sold according to the terms stipulated by the deploying organisation (Browne, 2015; Chun, 2005). With the spread of commercial applications, there is now a gap between the ethical vocabulary for AI (fairness, transparency, accountability and human dignity) and enforceable constraints in real-world systems (Mendoza, 2022; Ruster & Davis, 2025).

Many of the activities carried out by the companies in this field are known as “ethics washing,” which include voluntary, non-binding ethical declarations to avoid mandatory regulations (Mittelstadt, 2019; Wagner, 2019). However, this word is used loosely. The most prominent recent case in China, the dispute at Hangzhou Safari Park, does not have a published code of ethics. A commercial company unilaterally required pre-paid users to submit their Face Recognition Data or lose access in the absence of regulations. We believe that the two phenomena (published ethics statements and silent strategic under-investment) have the same micro-structure: both are cost-minimizing responses to a low expected sanction. Treating them as analytically independent fails to consider the relationship among them.

Most previous research on AI ethics has focused on how algorithms shape society. The former is not appropriate. A single case: how have the forces of social change, the media, the judiciary and the legislature altered the rules for facial recognition technology? In line with the first change in science and technology studies, this time we will also examine how society participates in the governance of artificial intelligence (Leese, 2022). The Hangzhou case is worth studying because it provides a typical case of this phenomenon: after a single consumer's refusal was reported by the media, a judicial interpretation was issued to support the law on biometric data.

The paper makes four contributions. Firstly, in concept, it extends the idea of ethics washing to include omission of investment, that is, strategic non-investment in the absence of an ethics charter, and it distinguishes an institutional responsibility gap from the traditional epistemic responsibility gap (Matthias, 2004). Secondly, theoretically, it provides a falsifiable responsibility-gap model and node-correlation and parameter-sensitivity analysis. Thirdly, in accordance with the theory-guided process tracing method, it records how a socially shaped governance path (litigation, media, judiciary and legislature) led to the closing of a regulatory vacuum, and it supplements the case with an exploratory fuzzy-set-inspired comparison of five governance regimes. Fourthly, in terms of methods, all quantitative results are supported by reproducible open-source datasets and scripts.

Section 2 reviews the previous research. Section 3 develops the model. Section 4 introduces the research methods, datasets and coding schemes. Section 5 shows the case history and the social forces that drove it. Section 6 is the model verification and exploratory contrast section. Section 7 introduces governance issues. Section 8 discusses external validity, global comparison, and limitations, and Section 9 concludes.

2. Literature and Research Gaps

2.1 *From Greenwashing to Ethics Washing*

Washing is a form of greenwashing; the company's environmental goals are inconsistent with its actions (Lyon & Maxwell, 2011; Schultz et al., 2025). The self-imposed rule of thumb has led to the regulation of AI in recent years (Dignum, 2022). Scholarly research on the ethics of artificial intelligence has produced three types of concepts. Compliance washing refers to the hiring of a third party to conduct an audit and taking no further action based on the results; discourse washing involves issuing ethical declarations without making any actual changes in practice; and strategic ethical commitments are used to delay the release of binding regulations. Wagner (2019) designated the last and most politically consequential type as: voluntary ethics commitments by European technology companies served as a way out of regulation. Mittelstadt (2019) found that the well-known principles of AI ethics are almost uniform in content and largely ineffective in practice.

Recently, AI ethics research has pointed out that many high-risk deployments are conducted without public ethics charters or compliance statements (Berridge & Grigorovich, 2022). However, this pattern of silent non-investment has not been included in the ethics-washing system. Most of the criticisms are that there are no words or actions, and they do not mention cases without words. We believe that the two patterns have the same structural logic and can be considered two types of ethics-washing.

Building on prior research, we extend these lines of inquiry to an underexplored area. Scholars studying greenwashing have categorized decoupling strategies into commissions (active symbolic claims) and omissions (silent non-disclosure or inaction) (Sidarta, 2026). On the other hand, much of the AI-ethics literature remains contested; companies have released charters, set up ethics committees, published principles, etc. Therefore, we have found two kinds of ethics-washing. Discursive ethics-washing is the proactive creation of symbolic commitments (charters, PR campaigns, ethics statements) to substitute for binding constraints and conform to the conventional definition (Biddle et al., 2024). Omission-based ethics-washing refers to the deliberate non-investment in a regulatory vacuum: the firm does not issue an ethics statement, fails to establish consent infrastructure, conducts no impact assessment, and neglects management of revocable data; thus, it avoids costs by exploiting consumer vulnerability and regulatory silence. Both types have the same structure: if the expected sanction is small, the firm will choose the least costly position of materiality in either form. The case of Hangzhou Safari Park is a typical case of the omitted sub-type, and therefore the classic concept cannot be applied to it. Analysis of the two types of intervention reveals that discursive washing is best countered by an audit of the gap between words and actions, and omission-based ethics-washing requires raising the baseline cost of non-investment itself.

2.2 *Biometrics, Power, and the Body*

Critical data studies and surveillance scholarship have added a dimension of power to the ethics-washing literature. Browne (2015) was worried that biometric surveillance would be employed to enforce racial and class control, and thus, the faces used by this technology would not be neutral.

Chun (2005) believes that biometric systems are redefining citizenship and Cath (2018) addresses this problem in the broader governance of artificial intelligence by trading bodily certainty for uncertain surveillance. Waelen (2023) has applied the above idea to current facial recognition and explored how biometric identification changes the boundaries of public and private life. The above research shows why forced biometric collection is not only a contract issue but also a reshuffling of the body-data-power relations. What they do not explain is how social actors respond; how to counter the power relations that have been exposed by the court, the media and legislation. This paper provides the solution.

Consequently, the notion of an ethical problem has also changed. Most AI ethics research has focused on bias and a lack of transparency in otherwise fair systems (Kellogg et al., 2020). From the perspective of critical theory, the problem is one of legitimacy; a face is taken because the deploying organisation holds the power to determine the terms of extraction, and thus, consumer consent is achieved under this asymmetry (Crawford, 2022; Pasquale, 2015). The Hangzhou case is a typical representative case: it was not a meeting of equals, but a unilateral modification of the park contract, and the choice was whether to show respect or lose the pre-paid access.

2.3 The Responsibility Gap and Regulatory Lag

The two senses of the responsibility gap are not the same and thus lead to different problems and solutions. The traditional epistemic responsibility gap is inherent in the technology (Matthias, 2004): learning automata whose behaviour is modified by experience are no longer predictable by their designers, and moral and legal responsibility presuppose knowledge and control that designers no longer possess. Sparrow (2007) advanced this idea for autonomous systems. The institutional responsibility gap, on the other hand, is external to the technology; it results from the delay in large-scale deployment and the non-entry into force of binding rules, so violations occur without a responsible party or penalty mechanism. The two defects can occur simultaneously; however, they are not the same. The Hangzhou case is mainly a problem of institutional gaps; it was not an unpredictable learning machine, but for about two years, there were no binding constraints on its operation.

The two deficiencies are frequently in combination. The epistemic gap and unpredictable system behaviour increase the institutional gap by making the design of regulatory rules more difficult, and delayed enforcement fails to address epistemic risks in practice. The main reasons for the institutional defects in the Hangzhou case were that the facial recognition system was technologically advanced and functionally predictable, and it was not an autonomous learning system operating independently of the designers' control. Thus, it is a relatively uncontaminated case of the institutional dynamics of regulatory lag.

The institutional gap has been enlarged by regulatory lag, that is to say, the delay between the introduction of new technology and its establishment as a binding constraint has made it difficult for rigid statutes to keep pace with agile technology (Webber, 1973; Marchant, 2011). However, the current studies have considered this deficiency to be a problem of philosophy rather than a state of

disequilibrium that social action can resolve. To address the above problem, a reform of the supervisory system is required, not just stronger regulation. Polycentric governance and co-production theory can be applied here.

2.4 Polycentric Governance and Co-Production

Although scholars have studied the actors of polycentric governance from all sides, in fact, it consists of three components (Ostrom, 2010): overlapping jurisdictions, autonomous rule-making at multiple scales, and collaborative learning across nodes. Recent work (Rittelmeyer et al., 2025) indicates that knowledge becomes governance only after institutionalisation of learning. For science and technology studies, Jasanoff (2004) proposes a theory of co-production, and governance order is not imposed on technology but jointly created through knowledge, law, discourse, and institutions. Polycentric governance, as shown in co-production, is not a regulatory system but a social process of jointly creating the rules for technology by courts, the media, the public and firms. The case in Hangzhou is a typical case of overlapping jurisdictions (contract, tort and personal information protection), autonomous rule-making by the Supreme People's Court, and collaborative learning (litigation informed the drafting context of the PIPL).

2.5 Research Gaps

Three gaps motivate this paper. Firstly, most of the ethics-washing literature focuses on discourse criticism; at present, there is a shortage of a micro-level model that explains the structural conditions favouring symbolic or absent investments, and it does not distinguish between published charters and silent under-investment. Secondly, most of the existing research on the problem of accountability is theoretical in nature; there is no model for the lag between implementation and enforcement, and how social actors help to address this gap has not been empirically studied. Thirdly, polycentric governance is only mentioned prescriptively; although the connections among litigation, the media and legislation are self-evident, there is little empirical evidence from outside the West on how co-production actually functions. The three are listed below.

The two concepts are the causes in a series. When the institutional responsibility gap is relatively large and the expected sanctions for non-compliance are weak, firms may be weakly motivated to comply with ethics or behave ethically due to cost reduction. Construct a micro-macro link model for structural regulatory conditions and corporate ethical behaviour in this paper.

3. A Responsibility-Gap Model

A model is an authoritative expression of falsifiable analysis that is neither predictive nor empirically verifiable. Formulate testable hypotheses on the impact of regulatory lag and oversight structure on corporate ethical investment. All numerical simulations below are illustrations of the theoretical logic, not statistical tests of a hypothesis. Qualitatively calibrate the parameters with the Hangzhou case in this paper to show the internal logic of the model; formal statistical tests will be performed in a large-N study later. It is not stated that an unobservable latent variable has been observed.

The two types of models for testing are: Proposition 1: *Ceteris paribus*, G_{peak} is positively associated with an increase in the length of the regulatory lag τ . Proposition 2: *Ceteris paribus*, polycentric oversight systems have a higher probability of detecting violations than monocentric systems, and the size of this difference decreases with increasing inter-node correlation ρ .

3.1 The Responsibility Gap as a Delayed Stock-Flow Process

We view the regulatory vacuum as a dynamic stock-flow system, and the accumulated unaddressed ethical exposure $E(t)$, the total stock of biometric privacy risk borne by the public that has not yet been offset by binding rules, increases over time as a result of the difference between incoming risk drivers and outgoing regulatory enforcement (Ahern, 2021).

The two reasons for the flow of ethical exposure are the intensity of technology rollout in commercial and public areas, denoted $T(t)$, and the severity of mandatory biometric data submission for end-users, denoted $S(t)$. The outflow of risk reduction is due to binding legal and regulatory enforcement, denoted $R(t-\tau)$, where τ is the regulatory lag: the time delay between the initial deployment of a technology and the entry into force of enforceable rules governing its use. The rate of increase in cumulative ethical exposure is shown by:

$$dE = \alpha T(t) + \beta S(t) - \gamma R(t-\tau) \quad (1)$$

where, α , β , and γ are weight coefficients capturing the relative contribution of each driver to overall ethical risk.

Integrating over time yields the total stock of cumulative unaddressed ethical exposure at any given point:

$$G(t) = \int_0^t [\alpha T(s) + \beta S(s) - \gamma R(s-\tau)] ds \quad (2)$$

At its base, this structure shows that the speed of new-technology introduction is slower than the pace of institutional reform. A larger regulatory lag τ results in an extended period of privacy risk for the public before the implementation of binding safeguards. We set $\tau \approx 25$ months to align with the time period of the Hangzhou case, from the park's unilateral introduction of mandatory face scanning in October 2019 to the full enforcement of China's Personal Information Protection Law (PIPL) in November 2021. The integration in Eq. (2) is well-defined for all $t \geq 0$; for $t < \tau$, the enforcement outflow is zero by construction, so $R(t) = 0$, as no enforceable regulatory instruments have entered into force yet. The cumulative exposure in the model is approximately one full lag period after the introduction of enforcement mechanisms, and then it gradually declines as regulatory outflows exceed new deployment inflows. Biometric data is inherently irreplaceably and irreversibly altered after collection, so damage caused during the period of regulatory delay cannot be fully rectified retroactively.

3.2 Strategic Under-Investment as a Signaling Game

The overall stock-flow dynamic above is the result of micro-level decisions by firms using biometric systems. A company that is not subject to external regulations has two options for ethics and different

costs and risks.

The first type of harm is strategic underinvestment (W), a minimal-compliance approach: the firm does not conduct formal privacy impact assessments, has no meaningful opt-out options, and fails to implement weak data governance measures; thus, it avoids penalties by relying on brand reputation, superficial user notifications, and the absence of active enforcement. This way has a low cost c_W .

The second post-design concept is ethics-by-design (D), and substantial ethical investment will be made in joint privacy impact assessments, revocable biometric templates, non-biometric access paths, and real-user choice. Therefore, the cost is higher and $c_D > c_W$.

Deployment benefits B accrue to the firm irrespective of the ethical position. The damage from under-investment is determined by how often the violation is found, q , and the size of the punishment after it is found, F . Therefore, the net payoffs of the two strategies are:

$$\pi_W = B - c_W - qF, \quad \pi_D = B - c_D \quad (3)$$

Strategic under-investment is a rational, cost-minimising option when its net benefit exceeds that of ethics-by-design, which reduces to the threshold condition:

$$c_D - c_W > qF \quad (4)$$

It is not economically viable, and there is also a power difference between the deploying organisation and the data subject. When a company has a diffuse base of individually weak claimants, such as pre-paid theme park customers who lack a collective action mechanism, the lower bound of the probability of accountability q approaches 0. Therefore, a relatively high fine will be ineffective in controlling the behaviour because the expected punishment is close to zero. Therefore, ethics washing and under-investment should not be considered moral failures of the individual firm but rather as predictable equilibrium results of a regulatory vacuum.

3.3 Polycentric Detection with Correlated Nodes

The single-level governance structure of China has only been strengthened at the provincial and municipal levels. The probability of finding a violation is thus the product of inspection frequency and sanction effectiveness:

$$q_{\text{mono}} = p_{\text{reg}} s_{\text{reg}} \quad (5)$$

where, p_{reg} is the probability of regulatory inspection and s_{reg} is the probability that inspection results in binding enforcement action. In the pre-PIPL context of the Hangzhou case, this monocentric detection probability was approximately 0.045, effectively negligible.

The many centres of power in the polycentric governance system are subject to public scrutiny and thus may be more likely to disclose information on a violation if such a disclosure occurs at any of these institutions. Under the standard assumption of independent node operation, the total detection probability is:

$$q_{\text{poly_indep}} = 1 - \prod_i (1 - p_i s_i) \quad (6)$$

Although mathematically convenient, the assumption of independence is often not realised in actual governance. Hangzhou's oversight nodes did not work independently but in a chain; the public attention they received led to changes in the judgments of the courts, which in turn prompted legislative modifications. Activation of one node reduces the threshold for activation of the next, and so on, forming a chain reaction.

To account for this correlation, we use a common-factor approximation of node correlation and employ the standard method of network reliability analysis. A single shared correlation coefficient $\rho \in [0,1]$ is introduced, and if $\rho=0$, then the nodes are completely independent; otherwise, $\rho > 0$. The actual detection rates of these nodes are:

$$r_{\text{eff},i} = r_i(1-\rho) + \rho \bar{r}, \quad (7)$$

where, $r_i = p_i s_i$ and $\bar{r} = \frac{1}{n} \sum_{i=1}^n r_i$. The total polycentric detection probability is then calculated as:

$$q_{\text{poly}}(\rho) = 1 - \prod_{i=1}^n (1 - r_{\text{eff},i})$$

with $r_i = p_i s_i$ and $\bar{r}$ their mean, and the total polycentric detection probability then follows as

$$q_{\text{poly}}(\rho) = 1 - \prod_{i=1}^n (1 - r_{\text{eff},i}), \text{ where } \rho \text{ belongs to the closed interval } [0, 1].$$

Two defects of this approximation are shown below: Firstly, the common-factor approximation can only find the degree of node synchronisation and fails to consider the sequential, order-dependent structure of real governance chains, such as litigation preceding media coverage and judicial reasoning, and so on. Second, as ρ approaches 1, the polycentric system degenerates into a single composite node, and the detection probability converges to that of the strongest node; however, this result is only applicable to the model and not representative of actual governance performance. Thus, the construction can be used as a teaching material. In this way, it keeps the main idea of polycentric governance theory: regulatory power is distributed among many parties in society instead of being concentrated in a single administrative department, and the correlation parameter shows that these nodes strengthen each other to form a cascade of social shaping rather than just an accumulation of individual efforts.

4. Methods

4.1 Case Selection: A Revelatory Single Case

We use an exploratory single-case design (Eisenhardt, 1989; Yin, 2018), analysed through theory-guided process tracing (George & Bennett, 2005). Process tracing investigates the causal mechanisms among events in a specific case. Since the purpose of this study is to demonstrate how the governance mechanism operates rather than to verify a causal law, it should be considered as mechanism illumination; the empirical materials provide rich details of the chronology and a reasonable explanation for the interaction among actors but do not establish a general causal relationship among them. Guo Bing vs. Hangzhou Safari Park is revelatory (Eisenhardt & Graebner,

2007; Siggelkow, 2007): it is China's first lawsuit involving facial recognition technology, and after about 18 months, an observable chain of events has been formed; at the same time, regulatory instruments in China have also been updated to close the regulatory gap. It is not typical of the industry; external validity will be presented in Section 8.

We adopt a theory-guided process-tracing design (George & Bennett, 2005) that tests whether the observable causal steps predicted by the responsibility-gap model are present in the case. We evaluate four typical features of process: temporality, sequence, mechanism and exclusion of other causes. The Goal is mechanism identification, not generalizable causal inference (Wang, 2026).

4.2 Data and Triangulation

We assembled a triangulated corpus of three channels. The first channel is legal texts: the two court judgments, the Civil Code (Standing Committee of the NPC 2020, effective January 1, 2021), the Supreme People's Court's provisions on facial-recognition civil cases (Supreme People's Court 2021; Fa Shi [2021] No. 15, effective August 1, 2021), and the Personal Information Protection Law (PIPL; Standing Committee of the NPC 2021), China's first comprehensive national legislation on personal information, effective November 1, 2021. The second is the corporate archive of yearly card contracts, store notices and SMS alerts. Third, by purpose selection, 35 media articles and expert opinions were chosen (Patton & Patton, 2002) from seven national and provincial outlets (including People's Daily, China News Weekly, and The Paper), covering the full dispute timeline from October 2019 to April 2021. Sampling prioritized analytical depth and range of perspectives, not statistical representativeness. Supplementary Table S1 shows the data corpus.

4.3 Coding and Data Structure

We conducted an inductive qualitative analysis informed by the Gioia methodology (Gioia et al., 2013). We developed first-order concepts from the source materials, grouped them into second-order themes, and organized the themes into aggregate dimensions. We used constant comparison during coding, drawing on procedures described by Corbin and Strauss (2008). The two coders independently conducted the coding, then discussed any differences in their results, and achieved a Cohen's kappa coefficient of 0.78 for the overlapping subset, comprising [number] of the [total number] coded units. Final coding was achieved by iterative reconciliation. Supplementary Table S2 shows the data structure.

5. Process Tracing: How Social Actors Closed the Gap

This section introduces the sequence of events from the individual's case to the media's coverage, judicial interpretations, and finally, legislative changes that addressed the regulatory gap lasting approximately 25 months, from the park's introduction of mandatory facial-recognition entry in October 2019 to the PIPL's entry into force on 1 November 2021 in the Hangzhou case. Table 1 is the entire event-behaviour-outcome sequence, and Figure 1 is shown below. Each stage corresponds to a different node in the polycentric governance model, and together they show how dispersed social actors

participate in the joint production of binding rules for biometric technology.

Table 1. Event-Behaviour-Outcome Chain of the Hangzhou Facial Recognition Dispute

Date	Event	Actor behaviour	Outcome
Apr 2019	Guo Bing purchases a two-person annual entry card (fingerprint + photo verification)	Consumer agrees to terms under the existing access regime	Baseline contractual relationship established
Oct 2019	Park unilaterally switches entry verification to mandatory facial recognition	Firm amends contract terms: Face Recognition Data submission required for access, with forfeiture of pre-paid entitlements for non-compliance	Mandatory biometric collection begins; S(t) shifts upward
Oct 2019	Guo Bing refuses mandatory face scanning and files suit at Fuyang District Court	Individual litigants initiate legal action independently; no class action mechanism available; litigation costs borne personally	Legal oversight channel activates; public litigant node enters operation
Nov 2020	First-instance judgment issued	Court awards 1,038 RMB compensation and orders deletion of facial biometric data; declines to invalidate park entry notices	Legal precedent established but remedy remains limited; regulatory outflow R(t) commences
Apr 2021	Second-instance judgment issued	Hangzhou Intermediate Court upholds first-instance ruling and expands order to include deletion of fingerprint data	Principle of biometric data destruction upon user request formalized
Aug 2021	Supreme People's Court issues judicial interpretation on facial	Judiciary establishes uniform national standards for biometric privacy	Overlapping jurisdictional governance takes effect; R(t) strengthens

	recognition civil cases (Fa Shi [2021] No. 15)	disputes	substantially
Nov 2021	Personal Information Protection Law enters into force	Legislature codifies separate consent requirements for sensitive biometric data, and necessity and proportionality principles	Binding national regulatory framework established; regulatory lag τ formally closes after approximately 25 months




Figure 1. The Co-Production Chain: From Individual Refusal to Legislative Codification

5.1 Firm Strategic Calculus and the Rationality of Under-Investment

Based on the cost-benefit analysis in the responsibility-gap model, the park has decided to require mandatory face scanning. Substantive ethics-by-design would have required to make substantial investments in several safeguards, such as providing a non-biometric entry route for users who opted out of facial scanning, creating fully revocable biometric templates, and establishing transparent data-handling protocols. In contrast, the actual strategic position of the park incurred a negligible cost: a single SMS notification and on-site notices; no public disclosure of technical specifications for data storage or use occurred.

In other words, there would be no penalty for such conduct. There was no specific national personal information protection law, no judicial interpretation of facial recognition technology, and the foundation of the park's pre-paid annual card users was scattered across different areas and lacked organisation for legal action. Given the above circumstances, the choice not to engage in substantive ethical investment can be viewed as a rational, cost-minimising response to the structural incentives of a regulatory vacuum, as shown by the threshold condition in Eq. (4), rather than negligence or a moral failure.

5.2 Individual Litigation as the Activating Node of Governance

Guo Bing, a university lecturer, was the first to initiate the construction of the governance chain. It was not a class action lawsuit; at that time, there was no collective litigation system for biometric privacy rights under Chinese law, and he had to bear all the financial and procedural costs of the lawsuit by himself.

Therefore, his decision to take up the case served as a form of citizen participation; an educated and well-resourced person was willing to assume asymmetric personal risks to act on behalf of all users who were unable to organise collectively. Therefore, according to the theory of polycentric governance, the public node for supervision is not uniformly accessible to all members of the public. It is necessary that a rare person arises who is willing and able to bear the disproportionate costs of individual legal action.

5.3 Media Amplification and the Construction of a Public Issue

The media's coverage has done more than just spread the news of a legal case. The scope of the dispute has expanded to a broader problem in the country: bodily integrity, the extent of state and commercial surveillance power, and the fundamental right to refuse biometric data collection; it is no longer a narrow private contractual dispute over a 1,038 yuan refund.

The discursive function is a characteristic feature of governance as presented in science and technology studies; that is, rules are not only formed by courts and legislatures but also negotiated in the public sphere through changes in shared understandings of acceptable behaviour.

Media coverage in our corpus operated in two ways. First, the commercial and investigative media platform The Paper used this case as the subject of research on biometric privacy and extended its coverage from a private contract dispute to the larger problem of bodily integrity. Later, the official organ of the central government, the People's Daily, issued a leading commentary and added the issue to the national policy list. Organize the 35 media articles and expert opinions in a two-step manner: first, report on an abnormal individual lawsuit; second, name it China's first legal test for facial recognition technology, and finally, present the dispute as an industry-wide warning signal. The first level of enlargement for a single person's lawsuit increased it to a regulatory reference case and thus prompted further judicial and legislative reforms.

5.4 Distributional Impacts and the Necessity Test

The costs of the responsibility gap are not spread evenly among the people. Compulsory biometric submission will be an extra burden for vulnerable groups; older adults unable to use the app-based alternative system will face inconvenience, children's biometric data will be collected without full informed consent, and low-income or rural users will find the prepaid annual card financially burdensome. None of these groups were part of the Hangzhou litigation, so they did not bear the consequences of regulatory failure.

In terms of welfare, the practice of biometric collection is ethically justified only if its marginal social benefit exceeds the expected harm to privacy, given the high sensitivity of bodily integrity:

$$\Delta U = B_{econ} - H_{priv} u(H) > 0 \quad (8)$$

where, B_{econ} is the economic efficiency gain, H_{priv} is the privacy harm, and $u(H) > 1$ is the welfare premium attached to violations of bodily data integrity.

The Hangzhou case showed only a small economic gain and limited the expansion of entry capacity. The loss of privacy was relatively high; as facial templates are permanent identifiers, only those who paid for access were subject to collection. According to the new legal provisions of both the Civil Code and the PIPL, the protection required for biometric data is relatively high. Therefore, the net welfare calculation is negative; it did not meet the necessity test, and at a low level of efficiency, a non-biometric option (fingerprint entry) was already available. The problem of harm to digitally marginalised groups is aggravated by this ethical lapse rather than reduced.

5.5 Evidence on the Transmission from Judgment to Legislation

Chronology does not prove that the ruling led to the new law. Therefore, we have gathered the three tiers of supporting evidence to document the paths of influence, but we have not exaggerated the results.

First, the judgment-to-interpretation path. The two main doctrinal rules in the second-instance judgment are that individuals have the right to request the deletion of their biometric data, and collection should meet a necessity test by using a non-biometric alternative when available. These judicial rationales later found formal expression in the subsequent Supreme People's Court judicial interpretation (Fa Shi [2021] No.15, effective August 2021). Article 12 of the Supreme People's Court's Provisions on Several Issues Concerning the Application of Law in the Trial of Civil Cases Involving the Processing of Personal Information by Using Facial Recognition Technology supports a request for deletion of facial information where the information processor has breached the parties' agreement. Article 4(1) provides that a processor generally may not condition a service on consent to facial-information processing unless that processing is necessary to provide the service. We have found that this is a correlation between judicial reasoning and the subsequent codification of doctrine; there is no mention of the case name, and we do not present a direct textual quotation.

Second, the interpretation-to-legislation channel. Article 28 of the PIPL classifies biometric information as sensitive personal information and permits its processing only for a specific purpose, where necessary, and subject to strict protective measures. Article 29 requires separate consent for processing sensitive personal information. Article 5 sets out the principles of lawfulness, legitimacy, necessity, and good faith. Scholarly reports on the drafting of PIPL consistently point out that the new judicial practice in biometric disputes informed the design of the sensitive-information system. As with the first channel, we will also consider this an indirect effect of the documentation.

Third, the media-to-agenda channel. National media, including state-owned media, reported on the case extensively, and at the same time, biometric privacy was added to the policy agenda through legislative proposals and representatives' motions for facial recognition. Corpus data shows the time lag between

media coverage and government attention to a problem, but it cannot be concluded that a particular news report caused a policy change.

To limit our causal claims, we examine the three competing explanations for the timing of biometric regulatory changes. First, the diffusion of international regulations: The world's discussions on GDPR and AI governance have provided a basis for China's digital regulations. Although it is a necessary condition, it does not explain why the specific judicial rules for biometrics were issued in the four months following the second-instance judgment. Secondly, there were some other concurrent cases of biometric privacy disputes in 2020 and 2021, but none attracted national media attention, were published as appellate judgments, or were cited in later judicial guidance. Third, central digital economy policy: Although there was a general trend to regulate platform economies, there was no specific mention of facial recognition. Together with the three channels above, we believe that the Hangzhou litigation was not the sole cause of the regulatory change; rather, it served as a crucial trigger and amplifier to accelerate the refinement of biometric-specific rules in the larger legislative process.

6. Model Calibration and Cross-Case Comparison

6.1 Model Calibration

The empirical governance timeline of the Hangzhou case is used to calibrate the model, and all parameter values are considered qualitative anchors rather than empirically obtained coefficients. Full calibration information is provided in Supplementary Table S3.

For state regulators, before the PIPL, the commercial biometric landscape had little regular inspection and weak contractual penalty clauses; thus, the base detection probability of p_{reg} was 0.15 and the sanction probability of s_{reg} was 0.30. Technical audits were not included in the Hangzhou case, so a hypothetical baseline detection probability of $p_{tech} = 0.20$ and a sanction probability of $s_{tech} = 0.50$ were set to reflect the general credibility of independent professional audits. At the time, individual legal action for public litigants concerning biometric privacy was extremely rare; thus, the detection probability p_{public} was 0.25, but the strong judicial recognition of Guo Bing's claim supports a higher sanction probability s_{public} of 0.60. Finally, the combined media and judiciary node has significant agenda-setting power and a strong precedential effect, and thus are worthy of detection ($p_{media} = 0.40$) and sanction ($s_{media} = 0.70$).




Figure 2. Conceptual Simulation of Cumulative Ethical Exposure $G(T)$

Parameters are calibrated to the case and are illustrative, not empirical measurements. Vertical dotted lines mark the Civil Code (January 2021), the Supreme Court facial-recognition interpretation (August 2021), and the PIPL (November 2021).

Figure 2 shows the simulated trajectory of cumulative ethical exposure $G(t)$ over the entire case period. In October 2019, after the park introduced mandatory face scanning, there was a sudden increase in cases and an increase in unregistered residents and non-compliant data submissions due to a regulatory gap. The three vertical dotted lines are successive enforcement stages, and at each stage, regulatory outflows are increased while the accumulation of risk is reduced. Based on the theoretical assumptions of the model, the cumulative exposure will reach a maximum after about one full lag period following the implementation of the first batch of binding enforcement instruments, and then gradually decrease as outflows exceed ongoing inflows. As shown in the model, regulatory lag is a structural reason for the accumulation of harm, and successive institutional responses gradually reduce the amount of unaddressed ethical exposure.




Figure 3. Polycentric Detection Probability across Node-Correlation Values

The monocentric baseline (state regulator only) is shown as a dashed line. Node probabilities p_i and sanction probabilities s_i are calibrated qualitatively to the pre-PIPL context (see text).

Figure 3 shows the total polycentric detection probability as a function of the degree of correlation among oversight nodes. A dashed line is a monocentric baseline, only state regulators are employed, and the detection probability is approximately 0.045. However, the four-node polycentric system has increased the detection rate of most correlation values significantly. At $\rho = 0.9$, the degree of correlation among the nodes is relatively high, and the polycentric detection is about 0.46, which is significantly higher than the monocentric baseline. This result is consistent with the main theoretical assumption of the model that polycentric supervision changes the firm's cost-benefit threshold in Eq. (4) independently of the correlation among nodes, although the size of this effect is reduced as the nodes become more aligned. As shown in the supplementary Table S4, the direction of the conclusion is unchanged for all example parameter values.

6.2 An Exploratory Fuzzy-Set-Inspired Comparison across Five Regimes

To determine whether the socially constructed governance mechanism in Hangzhou can be extended to other parts of China, we conducted an exploratory fuzzy-set-inspired comparative analysis of five national biometric governance systems: China (the Hangzhou case), the United States (Illinois Biometric Information Privacy Act, BIPA), the European Union (EU AI Act and General Data Protection Regulation (GDPR)), the United Kingdom (Metropolitan Police facial recognition trial), and India (Aadhaar face authentication).

Maximum-difference sampling is used to select the following five regimes: the three legal systems (civil law, common law, and mixed systems); the three stages of economic development; and the three types of governance models (administrative, judicial-deterrence, and socially driven). This Design can be used to explore institutional patterns in various places.

We calibrate four conditions: ex-ante legislative framework (A), private right of legal action (B), media and public activism (C), and independent technical audit (D). The result is the strength of enforceable biometric privacy protection. The membership scores of the fuzzy set range from 0 to 1, and they are assigned according to qualitative institutional anchors rather than survey data. Full calibration scores and their qualitative reasons are shown in Table 2, and these are also visualised in Figure 4. Given that there are only a few cases (N=5), we will not carry out the formal procedures of standard qualitative comparative analysis (such as truth-table minimization, consistency and coverage calculations, or necessity analysis). Fuzzy-set calibration is employed to offer an organised and transparent instrument for cross-case comparison; instead of verifying a hypothesis, it produces hypothesis-generating results.

Table 2. Fuzzy-Set Calibration of Five Biometric Governance Regimes (Membership 0 To 1)

Case	A: ex-ante legislation	B: private right of action	C: media/public activism	D: independent technical audit	OUT: enforceable protection
China (Hangzhou)	0.20	0.30	0.85	0.20	0.60
US (Illinois BIPA)	0.40	0.95	0.50	0.55	0.80
EU (AI Act/GDPR)	0.95	0.60	0.55	0.80	0.85
UK (Met FRT)	0.30	0.55	0.75	0.40	0.45
India (Aadhaar FRT)	0.15	0.25	0.60	0.10	0.20

Notes. Calibration anchors: EU A = 0.95 (AI Act ex-ante risk classification plus GDPR unified enforcement); US B = 0.95 (BIPA statutory damages plus private right of action); China C = 0.85 (sustained national media debate and public mobilization); UK C = 0.75 (significant public controversy over the Metropolitan Police facial recognition trial); India outcome = 0.20 (no dedicated biometric privacy statute). Scores are qualitative anchors, not survey estimates.




Figure 4. Fuzzy-Set Calibration: Condition Membership across Five Biometric Governance Regimes

Membership values are qualitative anchors, not survey estimates.

According to the standard fsQCA rules, first of all, a necessity test is performed for each condition; if the consistency score is 0.9 or higher, then that condition will be considered necessary to achieve the desired effect. The consistency scores for ex-ante legislation, private right of action, media and public activism, and independent technical audit are 0.71, 0.65, 0.70, and 0.62, respectively. None of the above conditions have been met, and based on our previous research, no single pillar of governance is adequate for providing strong biometric protection alone. Due to the small N, formal sufficiency analysis or truth-table minimization has not been carried out; therefore, the configuration solutions are underdetermined, and the analysis is exploratory and hypothesis-generating.

Two Patterns are proposed from this exploration. First, a single reason is not adequate for robust protection. China excels in media and public activism (C) but lags in ex-ante legislation (A) and independent technical audits (D), so it achieves only an intermediate result; the United Kingdom also scores well in public activism but performs poorly on ex-ante legislation and is thus insufficiently safeguarded. Therefore, activism alone is not sufficient to replace the guarantees provided by institutions.

Second, the two institutional models appear to have achieved good protection results in this small sample. First, introduce ex-ante legislation and an independent technical audit (as shown in D); second, build a risk-based AI Act system in the European Union and apply the GDPR via third-party conformity assessment. Secondly, powerful private rights of action have been created through enabling legislation (B and A), such as Illinois's BIPA system, which provides for statutory damages and allows

for private lawsuits to promote compliance. At the same time, the Chinese model lacks strong ex-ante legislation and independent audits; although there is public participation and judicial review in some respects, it is weaker than the Indian case, which performs poorly in all four aspects. We believe that the above alignments are in the exploratory stage and their purpose is to generate hypotheses for large-N studies, not to establish configurational sufficiency.

Together, the case study and the cross-regime comparison support the main idea of this paper: socially shaped governance offers a practical, albeit imperfect, alternative to the formal institutional system in the West for biometrics. However, it has not been widely used in practice; there is a need for a willing individual claimant and an amplifying public sphere.

7. Governance Implications: From Regulation to Social Shaping

The responsibility-gap model points out the intervention points that affect the probability of detection q and the extent of regulatory lag τ , rather than concentrating solely on the size of sanctions F . Each lever is connected to the mechanism shown in the Hangzhou case and interpreted through the lens of co-production; governance is not imposed on society from above, but rather a collaborative process built together with society.

7.1 The State: Metagovernance That Opens Space for Social Participation

The regulatory void at the time of the Hangzhou case spanned 25 months, and during this period, unconstrained deployment caused significant privacy harm. Instrumental tools can be introduced, such as regulatory sandboxes, algorithmic filing requirements and sunset clauses, to reduce the delay by allowing regulators to observe the deployment in a limited environment and modify rules before substantial harm occurs. The Government will also speed up new-regulation issuance. A more far-reaching function of state governance is to create institutional space for all parts of society—citizens, the media, independent auditors, etc.—to participate in and engage with the governance process. The Supreme People's Court's judicial interpretation on facial recognition, released just four months after the second-instance judgment in the Hangzhou case, is an example of this metagovernance function: it has reduced the effective regulatory lag by formalising legal standards in advance of legislative codification and embedded the substantive rights stated in Guo Bing's litigation into binding judicial doctrine. This sequence of judicial interpretation preceding legislative codification shows a typical pattern of Chinese regulatory governance for new-generation technologies. Instead of issuing detailed ex-ante statutes that are likely to be outdated quickly, authorities often allow for judicial practice and local experimentation to generate normative rules that are later incorporated into national laws. Hangzhou is a typical case of litigation-as-experimentation. Therefore, state institutions do not merely address deficiencies by themselves; rather, they create the circumstances for the transformation of social movements into stable norms.

7.2 Firms: Embed Social Values at the Front End

As shown in Eq. (4), increasing the cost of strategic under-investment c_w is as effective a deterrent as increasing the probability of detection q . Mantelero, 2022 proposed an obligatory algorithm impact assessment system for high-risk biometric applications, which would require enterprises to answer the following questions before deploying such systems: is there a non-biometric method of data collection available; can the biometric template be withdrawn by the user; what is the impact of deployment on various age groups and genders, and how will this be analysed. Furthermore, a simplified process for opting out of data collection should be implemented for users. The deployment of Hangzhou Safari Park did not meet any of the above indicators. By adding these requirements to binding procedural rules, the initial cost of under-investment is reduced before individual litigation leads to enforcement, and social values of bodily integrity and choice are embedded in the front-end technical design rather than being applied as an afterthought.

7.3 The Technical Community: Empowering Public Oversight

Opacity in systems and operational workflows generates unmanaged ethical hazards at their source. Technical stakeholders spanning researchers, data scientists and standards organizations play a role extending past ordinary firm audits: they break up knowledge monopolies, the core reason algorithmic black boxes escape accountability. Privacy-enhancing technologies including federated learning and differential privacy reduce mandatory data collection $S(t)$ by limiting raw data extraction. Independent third-party technical audits improve detection odds for violations that state regulators lack the manpower and expertise to identify. Public-facing documentation written in plain language for system design and data practices expands oversight beyond closed expert networks. The Hangzhou case lacked this component entirely. No technical audit evaluated the park's facial recognition system, meaning detection of its issues relied fully on a single private litigant's efforts.

7.4 The Public: From Data Subject To Governing Subject

The Hangzhou case was not driven into public view by coordinated regulatory action. It began with one litigant, gained traction through media attention, and gradually pushed governance issues onto the agenda. This trajectory indicates that public pressure can redefine an ethical problem even before formal regulation is in place. Institutionalizing that public role can occur through different channels, including collective litigation, public interest legal action, and community demands for oversight of biometric systems in public spaces. Giving affected people a meaningful voice in ethical review can improve detection more reliably than procedural compliance alone, because those who experience the harm are often best placed to identify it. At a deeper level, this approach recasts citizens not just as data subjects entitled to state protection, but as active participants in governing technology. The Hangzhou case provides empirical support for this view: public pressure was not peripheral to the outcome, but was the single factor that most helped close a 25-month regulatory gap.

Together, these four arguments point to a common conclusion. Ethical improvement should not be reduced to top-down correction or isolated technical fixes. It requires a polycentric governance system in which multiple actors jointly shape the rules of algorithmic systems. In short, the movement is from algorithms governing society toward society shaping algorithms.

8. Discussion

8.1 Theoretical Implications: Co-Production and Data Power

This study illustrates a dynamics of co-production: governance rules governing facial-recognition systems were not imposed upon firms externally; instead, they emerged from interactions among technical design, judicial reasoning, media debate, and legislative action (Jasanoff, 2004). This finding is also consistent with critical data studies, which emphasize that responsibility for data-related harm is not distributed evenly across actors, but tends to cluster around those who have the capacity to shape data systems and their broader social consequences. In this sense, the case supports arguments about dispersed responsibility and weak institutional safeguards (Browne, 2015). Polycentric governance is often presented as a more balanced alternative to top-down control, but it is not necessarily a neutral or purely technical solution. Its effects depend heavily on the capacity of different actors to mobilize resources, frame issues, and influence development pathways.

Beyond this, the Hangzhou case shows how co-production can operate at the level of technological design itself. The original facial recognition system, with its mandatory submission model and no meaningful opt-out option, was shaped not only by regulatory pressure but also by the company's desire to reduce operational friction. Litigants and civil-society groups did not simply add layers of oversight on top of an existing technology; they reshaped the terms under which the technology could be deployed. This is what co-production means in practice: social and technical change develop together, rather than being treated as separate processes. At the same time, this downward co-production is not necessarily empowering for everyone. It can shift power relations between firms and civil society, but it does not automatically fix deeper structural inequalities. The distributional effects remain uneven. Ethical review and technical governance are important, but they are not sufficient on their own. This is particularly clear when the most vulnerable groups—such as children, elderly people, rural communities, and low-income users—are disproportionately affected by biometric data extraction, yet lack the organized capacity to contest those practices effectively.

8.2 External Validity and Scope Conditions

The Hangzhou case is revelatory rather than statistically representative, so the responsibility-gap model should be generalized only under clearly specified scope conditions. The model applies most directly to high-stakes biometric deployments in commercial or public-service settings marked by three institutional features: technology moving faster than formal law, responsibility being dispersed across informal claims-making processes, and at least one independent actor having the institutional capacity to challenge non-compliant practices. The model is less relevant to low-risk administrative applications,

even where polycentric oversight structures formally exist. This aligns with the exploratory comparison in Section 6.2, which found that socially shaped governance mechanisms vary in strength across national contexts. The argument, therefore, is not that polycentric governance is universally effective. Its impact depends on institutional complementarities. A strong public role can improve oversight, but if that role remains partial and is not supported by other governance pillars, it will be constrained in delivering robust and enforceable biometric privacy standards.

8.3 Global Comparison: Three Pathways to Biometric Governance

The Chinese experience represents one of three distinct global pathways for biometric governance. In the United States, Illinois's Biometric Information Privacy Act (BIPA) establishes a framework built around private rights of action and statutory damages. This is a judicial-deterrence model: firms maintain compliance largely because they face substantial litigation risks. By contrast, the European Union leverages the AI Act to enforce ex-ante risk classification and conformity checks. This constitutes an administrative pre-commitment model, which embeds compliance requirements into products prior to their deployment. The Chinese pathway, illustrated by the Hangzhou case, is socially driven. A single lawsuit sparked judicial deliberation, whose conclusions were later formalised into legislation.

While these three pathways are shaped by different institutional resources and legal traditions, they rest on a shared core logic. Across all three systems, insufficient ethical investment by technology operators is challenged not by a single central authority, but by a network of diverse actors—including litigants, auditors, regulators, and members of the public. As a non-Western reference case, the Chinese example highlights how individual activism and media scrutiny can drive change in contexts where class-action lawsuits and independent technical auditing remain limited.

Each pathway is rooted in its own set of institutional foundations. The U.S. private-rights model draws on common law traditions, class-action tools and statutory-damages rules that reduce the barriers for private parties to enforce privacy rules. The EU's administrative model relies on a powerful central regulator and a well-established history of pre-market product risk oversight. The socially driven Chinese pathway emerges in a setting where collective litigation is constrained and independent technical audit bodies are immature. Here, public contestation paired with responsive courts becomes the main engine of governance change.

These pathways also come with their own inherent vulnerabilities. The U.S. private enforcement model fails to deliver consistent protection for claimants who cannot bear the costs of legal action, resulting in uneven, state-specific regulatory coverage. The EU administrative approach can be slow to adapt and disconnected from on-the-ground deployment. The Chinese socially driven model, meanwhile, hinges on two unpredictable prerequisites: a willing individual claimant, plus media channels capable of amplifying the dispute. Neither condition can be reliably guaranteed.

A truly robust global governance framework would integrate strengths from all three models: private rights of action, administrative standards that bind firms before deployment, and public discourse to shape regulatory priorities. For Global South jurisdictions with neither the US-style litigation infrastructure nor the EU's administrative capacity, the Chinese pathway carries both lessons and caveats. Social activism can help narrow regulatory lags, yet it remains a fragile replacement for institutional representation of digitally vulnerable groups. More broadly, the case suggests that public mobilization and judicial responsiveness can partially stand in for formal safeguards, but only as a second-best solution. Sustainable biometric privacy protection ultimately requires developing durable institutional pillars: proactive legislation, private enforcement rights, and independent audit systems that underpin stronger regulatory regimes.

8.4 Limitations and Future Research

This study has four key limitations.

First, the analysis cannot establish definitive causal claims. Although documentary records establish a clear sequence of events, we cannot draw strict causal conclusions. The drafting of the PIPL was already underway before the court handed down its ruling. This means the lawsuit accelerated the development of biometric data rules, rather than triggering the regulatory process from scratch. Broader public debates and international regulatory trends also act as unaccounted confounding variables. For these reasons, this work identifies a plausible causal mechanism instead of delivering conclusive causal evidence.

Second, the responsibility-gap model developed here is a context-specific conceptual framework, not a statistically generalizable one. Its parameters are tailored closely to the Hangzhou facial recognition dispute. Furthermore, the simplified representation of oversight nodes omits nuanced strategic behaviour, mixed corporate stances, and reputation-related interactions that shape firms' ethical decision-making in practice.

Third, the cross-regime comparison remains exploratory. Drawing on a small sample of five national governance systems, the analysis does not include formal truth-table tests or rigorous threshold validation. Future work should incorporate expert assessment to refine the qualitative scoring metrics used to compare different regulatory regimes.

Fourth, the case has limited representativeness. The litigation was brought by a well-resourced urban professional. Meanwhile, digitally vulnerable populations were not represented within this legal action, and those populations are disproportionately exposed to unregulated biometric surveillance. Future research should centre these marginalised user groups and explore inclusive governance arrangements that move beyond individual-led citizen activism.

9. Conclusion

Facial recognition brings to light a deep structural challenge within AI ethics. Where algorithmic systems remain opaque, data extraction follows commercial incentives, and regulation lags behind technological rollout, symbolic ethical commitments often stand in place of binding, enforceable safeguards. This paper formalizes this trade-off and frames the Hangzhou dispute as a concrete example of the resulting equilibrium. It then tracks how a socially assembled coalition helped narrow this regulatory delay, spanning an individual litigant, media outlets, courts and lawmakers. Conceptually, the work expands the notion of ethics washing to include passive underinvestment in risk mitigation. It also draws a clear distinction between this institutional responsibility gap and the better-known epistemic gap discussed in prior literature. Methodologically, the analysis draws on theory-guided process tracing alongside an exploratory comparison inspired by fuzzy-set approaches, supported by an open, reproducible dataset.

For practical policy, the takeaway is not simply to increase regulatory penalties. The priority is to strengthen the capacity of social actors to identify, challenge and halt harmful technology deployments, while ensuring that groups with the least ability to push back do not bear the worst burdens of this governance shortfall. Technology, at its core, ought to uphold human dignity and collective welfare. A polycentric governance model, in this sense, returns stewardship of technology to society. It seeks to sustain innovation without sacrificing bodily autonomy, privacy, and equal standing—the very foundations that make a digital society livable.

Author Contributions

Qiuting Zhang: Investigation, Data collection, Validation, Writing – Review & Editing. Ziyu Jia: Conceptualization, Methodology, Writing – Review & Editing, Supervision, Project administration. Yuanyuan Li: Conceptualization, Methodology, Writing – Original Draft, Formal analysis.

Funding

Shandong Provincial Social Science Planning Project (Youth Project), Research on Mechanism Innovation and Policy Coordination of Industrial Chain and Innovation Chain Integration of Shandong's "Specialized, Refined, Differential and Innovative" Enterprises in the Digital Intelligence Era (Grant No. 26DGLJ16).

Institutional Review Board Statement

Ethical approval was not required for this study. The study did not involve human participants, human biological materials, interventions, or the collection of personal/identifiable information (e.g., names, contact details, images, recordings, survey responses, or interview data). Therefore, research ethics committee review was not applicable for the present analysis.

Informed Consent Statement

Not applicable.

Data Availability Statement

The data that support the findings of this study are available from the corresponding author upon reasonable request.

Acknowledgments

The authors thank the open-access judicial judgments and publicly released media corpus that made this case analysis possible. We also appreciate helpful comments from seminar participants on earlier drafts of this paper

Conflicts of Interest: The authors declare no conflict of interest.

References

- Ahern, D. (2021). Regulatory Lag, Regulatory Friction and Regulatory Transition as FinTech Disenablers: Calibrating an EU Response to the Regulatory Sandbox Phenomenon. *European Business Organization Law Review*, 22(3), 395-432. <https://doi.org/10.1007/s40804-021-00217-z>
- Ananny, M., & Crawford, K. (2018). Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society*, 20(3), 973-989. <https://doi.org/10.1177/1461444816676645>
- Berridge, C., & Grigorovich, A. (2022). Algorithmic harms and digital ageism in the use of surveillance technologies in nursing homes. *Frontiers in Sociology*, 7, 957246. <https://doi.org/10.3389/fsoc.2022.957246>
- Biddle, J. B., Nelson, J. P., & Olugbade, O. E. (2024). How Can We Know if You are Serious? Ethics Washing, Symbolic Ethics Offices, and the Responsible Design of AI Systems. *Canadian Journal of Philosophy*, 54(4), 329-345. <https://doi.org/10.1017/can.2025.9>
- Browne, S. (2015). *Dark Matters: On the Surveillance of Blackness* (p. dup; 9780822375302/1). Duke University Press. <https://doi.org/10.1215/9780822375302>
- Cath, C. (2018). Governing artificial intelligence: Ethical, legal and technical opportunities and challenges. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2133), 20180080. <https://doi.org/10.1098/rsta.2018.0080>
- Chen, W. (2026). Don't kick my ladder: How China transduces de facto RRI in the governance of facial recognition technology. *Journal of Responsible Innovation*, 13(1), 2611500. <https://doi.org/10.1080/23299460.2025.2611500>
- Chun, W. H. K. (2005). *Control and Freedom: Power and Paranoia in the Age of Fiber Optics*. The MIT Press. <https://doi.org/10.7551/mitpress/2150.001.0001>

- Corbin, J., & Strauss, A. (2008). *Basics of Qualitative Research (3rd ed.): Techniques and Procedures for Developing Grounded Theory*. SAGE Publications, Inc. <https://doi.org/10.4135/9781452230153>
- Crawford, K. (2022). Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence. *Perspectives on Science and Christian Faith*, 74(1), 61-62. <https://doi.org/10.56315/PSCF3-22Crawford>
- Denzin, N. K. (2009). The elephant in the living room: Or extending the conversation about the politics of evidence. *Qualitative Research*, 9(2), 139-160. <https://doi.org/10.1177/1468794108098034>
- Dignum, V. (2022). *Responsible Artificial Intelligence—From Principles to Practice* (Version 1). arXiv. <https://doi.org/10.1145/3582524.3582529>
- Eisenhardt, K. M. (1989). Building Theories from Case Study Research. *The Academy of Management Review*, 14(4), 532. <https://doi.org/10.2307/258557>
- Eisenhardt, K. M., & Graebner, M. E. (2007). Theory Building From Cases: Opportunities And Challenges. *Academy of Management Journal*, 50(1), 25-32. <https://doi.org/10.5465/amj.2007.24160888>
- George, A. L., & Bennett, A. (2005). *Case studies and theory development in the social sciences*. MIT Press.
- Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013). Seeking Qualitative Rigor in Inductive Research: Notes on the Gioia Methodology. *Organizational Research Methods*, 16(1), 15-31. <https://doi.org/10.1177/1094428112452151>
- Jasanoff, S. (Ed.). (2004). *States of Knowledge* (0 ed.). Routledge. <https://doi.org/10.4324/9780203413845>
- Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at Work: The New Contested Terrain of Control. *Academy of Management Annals*, 14(1), 366-410. <https://doi.org/10.5465/annals.2018.0174>
- Leese, M. (2022). Fixing State Vision: Interoperability, Biometrics, and Identity Management in the EU. *Geopolitics*, 27(1), 113-133. <https://doi.org/10.1080/14650045.2020.1830764>
- Lyon, T. P., & Maxwell, J. W. (2011). Greenwash: Corporate Environmental Disclosure under Threat of Audit. *Journal of Economics & Management Strategy*, 20(1), 3-41. <https://doi.org/10.1111/j.1530-9134.2010.00282.x>
- Mantelero, A. (2022). *Beyond Data: Human Rights, Ethical and Social Impact Assessment in AI* (Vol. 36). T.M.C. Asser Press. <https://doi.org/10.1007/978-94-6265-531-7>
- Marchant, G. E. (2011). Addressing the Pacing Problem. In G. E. Marchant, B. R. Allenby, & J. R. Herkert (Eds.), *The Growing Gap Between Emerging Technologies and Legal-Ethical Oversight* (Vol. 7, pp. 199-205). Springer Netherlands. https://doi.org/10.1007/978-94-007-1356-7_13

- Matthias, A. (2004). The responsibility gap: Ascribing responsibility for the actions of learning automata. *Ethics and Information Technology*, 6(3), 175-183. <https://doi.org/10.1007/s10676-004-3422-1>
- Mendoza, C. (2022). The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power: by Shoshana Zuboff, New York, Public Affairs, 2019, 704 pp., €20.00, ISBN 978-1610395694. *Church, Communication and Culture*, 7(2), 452-455. <https://doi.org/10.1080/23753234.2022.2086891>
- Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence*, 1(11), 501-507. <https://doi.org/10.1038/s42256-019-0114-4>
- Ostrom, E. (2010). Polycentric systems for coping with collective action and global environmental change. *Global Environmental Change*, 20(4), 550-557. <https://doi.org/10.1016/j.gloenvcha.2010.07.004>
- Pasquale, F. (2015). *The Black Box Society: The Secret Algorithms That Control Money and Information*. Harvard University Press. <https://doi.org/10.4159/harvard.9780674736061>
- Patton, M. Q., & Patton, M. Q. (2002). *Qualitative research and evaluation methods* (3 ed.). Sage Publications.
- Rittel, H. W. J., & Webber, M. M. (1973). Dilemmas in a general theory of planning. *Policy Sciences*, 4(2), 155-169. <https://doi.org/10.1007/BF01405730>
- Rittelmeyer, P., Lubell, M., Hovis, M., Heikkila, T., Gerlak, A., & Pozzi, T. (2025). Knowledge is not power: Learning in polycentric governance systems. *Review of Policy Research*, 42(3), 444-465. <https://doi.org/10.1111/ropr.12606>
- Ruster, L. P., & Davis, J. L. (2025). The Gaps that Never Were: Reconsidering Responsible AI's Principle-Practice Problem. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency* (pp. 350-360). <https://doi.org/10.1145/3715275.3732024>
- Schultz, M. D., Conti, L. G., & Seele, P. (2025). Digital ethicswashing: A systematic review and a process-perception-outcome framework. *AI and Ethics*, 5(2), 805-818. <https://doi.org/10.1007/s43681-024-00430-9>
- Sidarta, A. L. (2026). *Between Overclaiming and Silence: A Systematic Literature Review and Normative Ethics Framework for Corporate Sustainability Disclosure*.
- Siggelkow, N. (2007). Persuasion With Case Studies. *Academy of Management Journal*, 50(1), 20-24. <https://doi.org/10.5465/amj.2007.24160882>
- Sparrow, R. (2007). Killer Robots. *Journal of Applied Philosophy*, 24(1), 62-77. <https://doi.org/10.1111/j.1468-5930.2007.00346.x>
- Standing Committee of the National People's Congress of China. 2020. Civil Code of the People's Republic of China; Effective 1 January 2021. NPC: Beijing, China.

- Standing Committee of the National People's Congress of China. 2021. Personal Information Protection Law of the People's Republic of China; Effective 1 November 2021. NPC: Beijing, China.
- Supreme People's Court of China. 2021. Provisions on Several Issues Concerning the Application of Law in Trying Civil Cases Related to Facial Recognition Technology (Fa Shi [2021] No. 15); Effective 1 August 2021. SPC: Beijing, China
- Waelen, R. A. (2023). The struggle for recognition in the age of facial recognition technology. *AI and Ethics*, 3(1), 215-222. <https://doi.org/10.1007/s43681-022-00146-8>
- Wagner, B. (2019). Ethics As An Escape From Regulation. From "Ethics-Washing" To Ethics-Shopping? In E. Bayamlioglu, I. Baraliuc, L. A. W. Janssens, & M. Hildebrandt (Eds.), *BEING PROFILED* (pp. 84-89). Amsterdam University Press. <https://doi.org/10.1515/9789048550180-016>
- Wang, Y. (2026). Process Tracing and the Problem of Evaluating Causal Relationships Within a Case. *Philosophy of the Social Sciences*, 56(4), 311-346. <https://doi.org/10.1177/00483931261441189>
- Yin, R. K. (2018). *Case study research and applications: Design and methods* (Sixth edition). SAGE.

Supplementary material

The following tables support the main text. All data and code are reproducible from the provided scripts.

Table S1. Data Corpus

Data category	Specific sources	Period	Purpose
Legal and policy	First-instance judgment (Fuyang, 20 Nov 2020); second-instance judgment (Hangzhou Intermediate, 9 Apr 2021); Civil Code; PIPL; Fa Shi [2021] No. 15	2019-2021	Boundaries and lag measurement
Corporate archives	Annual-card agreement; store notices; SMS; unilateral contract amendment	Apr-Oct 2019	Mandatory-submission mechanism
Public discourse	35 news reports and	Oct 2019-Apr 2021	Media-mediated

expert commentaries
from 7 national and
provincial outlets

enforcement and
agenda setting

Table S2. Gioia Data Structure

Aggregate dimension	Second-order theme	Representative first-order concept
Intrinsic technological attributes	Biometric data vulnerability	Face cannot be reset; template irreversibility; no encryption disclosure
	Algorithmic opacity	No technical protocol disclosed; third-party vendor unknown; consent unintelligible
Socioeconomic drivers	Capital-driven efficiency overreach	Labour cost at gates; throughput logic; unilateral contract amendment
	Information and power asymmetry	All-or-nothing opt-in; pre-paid access held hostage; no real alternative
Governance-system failure	Ex-ante regulatory vacuum	No PIPL yet; no necessity or proportionality standard; only contract law available
	Reactive, media-triggered oversight	Enforcement only after media outcry; asymmetric litigation cost

Table S3. Parameter Calibration

Parameter	Value	Justification
alpha (deployment weight)	0.60	Modest exposure per unit commercial deployment
beta (submission weight)	1.00	Mandatory biometric submission is the highest-weight inflow
gamma (enforcement weight)	0.90	A binding instrument discharges exposure near one-to-one
tau (regulatory lag)	25 months	Oct 2019 mandatory submission

T(t) shape	0 before 2019-10; 1.0 to 2021-11; 0.2 thereafter	to Nov 2021 PIPL Deployment continues but constrained post-PIPL
S(t) shape	1.0 from 2019-10 to 2021-04; 0 thereafter	Mandatory submission stopped after deletion order

Table S4. One-At-A-Time Sensitivity of Peak Cumulative Ethical Exposure

Parameter	Tested values	G_peak range
α	0.30 / 0.60 / 0.90	25.5 to 40.5
β	0.60 / 1.00 / 1.40	25.8 to 40.2
γ	0.60 / 0.90 / 1.20	33.0 (stable in this window)
τ	12 / 24 / 36 months	30.8 to 34.4

Note. Values are produced by repro_analysis.py.