

The Publication Projection

Public Utility Without Public Access

Flyxion

Independent Researcher

October 2026

Abstract

Debates over public data usually treat the choice as one between releasing an observation and withholding it. This essay argues that the more useful object of analysis is the publication projection, the map from what an instrument observes to what an institution is authorized to release. Public utility, on this view, does not require public access to the underlying observation. It requires a defensible transformation from observations into outputs sufficient for a specified public purpose, and any claim of public-domain status attaches to those outputs rather than to the apparatus or its raw take, with authorization to release an output kept distinct from the legal rights that govern its later use. Roadside hazard reporting for connected vehicles serves as the principal worked example, radiographic quality assessment as a secondary case, and genomic statistics as the limiting case in which the principle strains. The essay states a constrained information-minimization model, shows that the model by itself settles none of the questions that matter, and separates four properties, output minimization, non-retention of source observations, non-acquisition of identifying information, and verifiable enforcement, none of which entails the others. It does not claim that the principle resolves ownership, consent, or surveillance. It closes on a question it leaves open: whether an institution can legitimately claim to operate a publication-only instrument while retaining the technical capability to turn that instrument into a surveillance device.

1 Observation and Publication

A sensor observes more than anyone needs to know. A camera mounted above a highway records every vehicle, every plate, every face behind a windshield, and the condition of the pavement besides. A radiograph records a skeleton, a set of implants, a date, and sometimes a name burned into a corner of the film. A sequencer records the whole inherited text of a person and, implicitly, a portion of the texts of their relatives. In each case the instrument

is capable of far more than any public purpose requires, and the usual policy conversation proceeds as though there were only two settings: the data are released, with whatever risks that carries, or they are withheld, with whatever benefits that forfeits.

That framing hides the object that actually does the work. Between the observation and the world sits a transformation, and the transformation determines what leaves the instrument and enters circulation. This essay calls it the publication projection. It may be as crude as a crop and a redaction, or as elaborate as a trained classifier that reduces a video stream to a short list of typed events. What matters is that it is a distinct component, with a definable input, a definable output, and a definable owner, and that the question of what belongs in the public domain is a question about its outputs.

The thesis is that public utility does not require public access to the underlying observation. A road user benefits from knowing that debris lies in the left lane two kilometers ahead. That benefit does not depend on access to the footage from which the fact was derived, and it is not enlarged by such access. The public interest in the fact is real and the public interest in the footage is, for this purpose, nil. The same structure recurs wherever an observation is collected by or on behalf of the public and then used for a narrow purpose, and it suggests that the public domain should be defined by what a system is authorized to release, not by everything its instruments are capable of observing.

Four clarifications limit the scope of the claim. First, the thesis concerns the content of releases and says nothing, by itself, about who may collect, who may retain, or who may inspect. Second, it does not assert that any particular projection is safe. A projection can fail in ways that are properties of the data and of the world around the data, and later sections describe several. Third, authorization to release an output is not the same thing as public-domain status. Authorization determines whether an institution may release a particular output, while public-domain status concerns the legal rights that govern its subsequent use. A projection can establish a technical boundary between what is observed and what is released, and an institution can adopt a rule that releases the outputs beyond that boundary without restriction, but neither step by itself makes the outputs public domain in a legal sense, which depends on law and on who holds rights in them. The essay uses the phrase public domain for the policy aim of unrestricted reuse of authorized outputs and makes no claim about existing legal entitlement. Fourth, the argument is a proposal and not a proof. The essay offers reasons to think the architecture is worth building and identifies conditions under which it would fail, and it declines to claim more than that.

2 Roadside Hazard Reporting

Networks of roadside cameras have been deployed in large numbers for the purpose of reading license plates. The same hardware, pointed at the same road, can serve a different purpose: detecting obstacles and reporting them to vehicles so that they know the conditions ahead. A connected vehicle approaching a curve benefits from learning that a stopped vehicle, a fallen load, a patch of ice, standing water, or an animal on the shoulder lies beyond its own sensing range. Vehicle-to-everything communication already provides the channel through which such reports would travel. What is missing is a settled account of what the roadside node is permitted to say.

The proposal is that the authorized public output of such a node be a short event record and nothing else. A record names a hazard class, an approximate location, a timestamp, optionally a confidence value, and an expiry time after which the report is withdrawn. It does not name the vehicle that stopped, the person driving it, the plate, the occupants, or the property adjacent to the roadway. The public system does not need these facts to function and, under the proposal, never receives them. Releasing such records into the public domain means that anyone may receive, store, redistribute, and build upon the hazard feed without license, so that a navigation service, a municipal maintenance crew, an independent researcher, and an automaker all draw from the same open stream.

The reasons for placing the feed in the public domain are straightforward. A hazard is a fact about a shared road and the road is public infrastructure. Enclosing the fact in a proprietary feed fragments the safety information, so that a driver learns of a hazard only if the vehicle belongs to the right ecosystem. Open release also permits independent verification: a record that anyone may collect can be compared against what anyone sees, and a node that reports hazards that are not there, or fails to report ones that are, can be identified by outsiders.

The design earns its strength from where the projection runs. If the node computes the event record at the sensor and transmits only that, then identifying imagery never leaves the device, and the privacy property is a consequence of the architecture rather than of a policy applied afterward. This is a stronger position than anonymizing imagery downstream, because a scrub performed downstream presupposes that the unscrubbed material was transmitted, stored, and available to whoever held it at the time.

The projection nevertheless leaks, and the leaks are worth stating plainly. A hazard event on a road traveled by few vehicles discloses more than the same event on a busy highway, because the set of persons who could have caused or experienced it is small. A stopped-vehicle report on a private lane with a single household at its end identifies that household in practice, whatever the record omits. A candidate safeguard is an aggregation

or delay rule that depends on traffic density, so that reports from sparsely used segments are withheld, coarsened in location, or released after an interval long enough to separate them from the events that produced them. Such a rule is a reasonable component of a projection and it should not be presented as sufficient for anonymity. A sparsely traveled road can remain identifying after substantial delay, since the inference rests on how few persons use the road and not on how recent the report is. Repeated releases compound the problem: a sequence of individually innocuous records, differenced against one another or against an auxiliary source such as a posted schedule or a public calendar, can permit reconstruction that no single record would allow. The density rule reduces the risk and does not remove it, and the essay claims no more for it.

The rule also collides with the purpose that justifies the system. A stopped-vehicle warning is valuable only while an approaching driver can still act on it, and a report held back until it can no longer be linked to a household may arrive after the obstruction has been passed. In that case the safeguard has not merely reduced utility, it has defeated the purpose, since a warning delivered too late is no warning. Suppression or delay therefore cannot be applied as a uniform default. It has to be weighed against the severity of the hazard, the reaction time a driver or vehicle actually has, and whether the same protection can be obtained another way, for instance by coarsening the reported location instead of delaying it, by directing a live report only to vehicles within a short range while the hazard is active and publishing a delayed, aggregated record afterward, or by withholding reports from segments where no approaching traffic could use them. Which hazards justify immediate disclosure, and how much residual risk to the residents of a sparse road is acceptable in exchange, are decisions about whose interests prevail when they conflict, and no information-theoretic analysis supplies the answer. This is the first concrete place where the utility threshold of the next section shows itself to be a governance choice.

3 A Constrained Formulation and Its Limits

The structure of the proposal can be stated formally, with the understanding that the formalism makes the structure inspectable and does not make it correct. Let $\mathcal{X}$ denote the observation space of the sensor and let X be a random observation drawn from it. Let $\mathcal{Y}$ denote the authorized public output space. A projection that may be randomized is not an ordinary function from $\mathcal{X}$ to $\mathcal{Y}$. It is a Markov kernel $P(y \mid x)$, which assigns to each observation a probability distribution over outputs, and a deterministic projection is the special case in which every such distribution is concentrated on a single output. Write Y for the released output, drawn from $P(\cdot \mid X)$. Let S denote the sensitive information associated with the observation, such as the identity of the vehicle or its occupants. Utility

is taken to be expected task performance and not a property of any single output. Let T denote the state of the world that the public purpose concerns, such as the presence and class of a hazard, and let $\ell(T, Y)$ be a loss that prices a missed hazard, a false warning, and a late warning at their different costs, so that the asymmetry between failures is part of the definition and not an afterthought. Then $U(P) = -\mathbb{E}[\ell(T, Y)]$, with the expectation taken over observations and, where timeliness matters, over the stream of reports. The design problem is

$$\min_P I(Y; S) \quad \text{subject to} \quad U(P) \geq u_{\min},$$

where I denotes mutual information. This is a privacy-funnel style formulation of the kind studied in the information-theoretic privacy literature [4], and it is related to the information bottleneck [7].

The first observation about this problem is that it is degenerate without its constraint. A constant projection, one that emits the same output regardless of the observation, drives $I(Y; S)$ to zero and discloses nothing. The unconstrained problem therefore has a trivial solution, and every nontrivial projection is justified only by the utility constraint. It follows that the choice of $u_{\min}$, the definition of U with the losses inside it, and the authority to revise any of these are not mathematical consequences of the model. They are institutional decisions about which public purposes are worth the residual disclosure, who defines them, and who may change them later. The mathematics can make those decisions inspectable, by forcing them to be stated as a threshold and a measure that others can examine and dispute. It cannot make them legitimate, because legitimacy depends on who holds the authority to set the threshold and on what recourse exists against that authority.

The second observation is that the formulation is an idealization and not a guarantee of privacy. Mutual information is an average-case quantity, and rare cases are exactly where harm concentrates, since a low average disclosure is compatible with a severe disclosure for a few individuals. The quantity as written also ignores auxiliary information. If an adversary holds a side channel Z correlated with S , the relevant quantity is closer to $I(Y; S \mid Z)$, which may be considerably larger than $I(Y; S)$, a point made forcefully by work on re-identification from supposedly anonymized releases [6, 5]. Finally, the formulation treats a single release, whereas a deployed system emits a stream, and the information in a sequence of outputs can exceed what any one output reveals. Worst-case and composition-aware frameworks such as differential privacy [1] address some of these concerns at a cost in utility, and a serious deployment would need to choose among them. The present formulation is best understood as a way of stating what must be decided and not as the decision.

Four properties of a publication regime are easily conflated and should be kept apart.

Definition 1. *Output minimization* is the property that the released output carries little

information about the sensitive attributes beyond what the specified purpose requires. *Non-retention* is the property that source observations are not persisted beyond the time needed to compute the projection. *Non-acquisition* is the property that the instrument is physically incapable of capturing the identifying information in the first place. *Verifiable enforcement* is the property that an independent party can confirm that the deployed device computes the stated projection and persists nothing else.

Non-acquisition differs in kind from the other three. Output minimization, non-retention, and verifiable enforcement describe how a particular configuration behaves, and each can change with a software update. Non-acquisition is a hardware-level constraint on what the apparatus can sense at all, and it is the only one of the four that does not depend on the operator’s continued compliance.

None of these entails the others. A device can emit only terse hazard records while writing all of its raw footage to local storage, and so satisfy output minimization while failing non-retention. A device can retain nothing while streaming rich, identifying output, and so satisfy non-retention while failing output minimization. A camera can momentarily hold identifying imagery in memory while computing a hazard record and then discard it, so that it satisfies non-retention without satisfying non-acquisition, whereas a sensor whose optics and resolution cannot resolve a plate or a face satisfies non-acquisition regardless of what its software does. A device can satisfy the first three properties and yet be impossible for an outsider to confirm as doing so, because its firmware is closed and its behavior under a changed configuration is unobservable, in which case verifiable enforcement fails. That property is also the hardest, since it is a claim about counterfactual behavior, about what the device would do if it were instructed differently, and an attestation of its current state establishes only that the state was what it was when the attestation was made. Two sensors that emit identical public outputs can thus differ in whether they ever held identifying information and in whether they could, and those differences leave the outputs unchanged while altering the security properties and the governance each arrangement requires.

4 Radiographic Quality Assessment

Medical radiographs present a case in which the public interest in the data is strong and the identifying content is concentrated in places that can be addressed. The research and educational value of large open collections of images is considerable, and the images themselves were produced by procedures that patients generally had little practical choice about. The usual identifying content, such as patient names, dates, record numbers, and device identifiers, resides largely in metadata and in text burned into the image, and it can be removed by well-understood procedures. The remaining content is not metadata at all.

Surgical hardware can carry serial numbers. Dental work and the morphology of particular bones can function as biometric features. A volumetric scan of the head can be rendered into a recognizable face. A projection that treats a chest film and a full cranial volume under the same rule is therefore a poor projection.

The role of quality assessment is to supply the information that selects among projections. A score computed over a radiograph, reflecting exposure, positioning, the presence of artifacts, and resolution, measures the diagnostic and research value of the image. A second assessment, of anonymizability, measures how much identifying structure survives the standard scrub for the modality and anatomy in question. The two are separate objectives and they pull against each other, since the properties that make an image diagnostically valuable, fine resolution, complete anatomical coverage, and clear bony detail, are largely the properties that make it identifying, so that improving quality can raise research value and disclosure risk together. A high score on the first therefore says nothing reassuring about the second. Together they determine the treatment of each image. A film of high research value and high anonymizability might be released after scrubbing. A film of high value and low anonymizability might be released only as a derived assessment, such as a quality grade or a set of measured features, while the image itself remains in a governed clinical setting. A film of low value might not be released at all, since the residual disclosure buys nothing. The projection here is a family of maps indexed by the assessment, and the assessment is itself a published quantity with its own utility and its own leakage.

The essay does not assume that this arrangement is free of difficulty. The anonymizability assessment must be reliable, and an assessment that fails to recognize an identifying implant defeats the scheme at the point where it matters. The family of projections is only as good as the classifier that routes images among them, and the classifier is a component whose errors have asymmetric costs. Even a reliable anonymizability assessment is an estimate of residual identifiability and not a guarantee of privacy, and it should not be cited as one. The case nonetheless illustrates the thesis cleanly. What enters the public domain is a transformation of the observation chosen for its fitness to a purpose, while the observation itself remains under whatever governance the clinical setting provides.

5 Genomic Statistics as the Limiting Case

The genome is the case in which the principle is least comfortable, and it is included because a principle that has not been tested against its hardest instance has not been tested. A genome is not a record that contains an identifier. It is, in the relevant sense, the identifier, and it additionally discloses information about relatives who never took part in any decision to release it. The tools that worked for hazard events and for radiographs, removing a

name or reducing an image to a score, have no purchase on a sequence that identifies its owner by construction. Surname inference from genetic markers [2] and the detection of an individual’s presence in a pooled summary [3] both show that the apparent boundary between individual data and aggregate data is weaker for genomes than for most other observations.

The candidate projection is the release of derived population-level statistics in place of individual sequences, such as allele frequencies or association summaries over a sufficiently large cohort. This preserves much of the research utility, and it fails in the places that the literature has documented. Aggregates drawn from small populations or from variants of high informativeness permit membership inference. Repeated releases over overlapping cohorts permit reconstruction. Linkage with other datasets permits re-identification. In each case the guarantee that the projection offers is weaker than in the hazard case, because the sensitive information is not a separable attribute of the scene that the projection can decline to capture. It is entangled with the very signal that gives the data its value.

For this reason the essay does not claim that individual genomes belong in the public domain under the publication-projection principle, and it regards that claim as unsupported. The principle supports, at most, the release of carefully chosen aggregates, with the stated understanding that the privacy claim is substantially weaker than for the first two cases, and that the consent of relatives who are affected without having been asked remains a problem that no projection addresses. The genomic case therefore marks the boundary of the argument. The three examples lie along an axis from a setting in which identifying content can be discarded at the sensor, through one in which it can be addressed by a routed family of scrubs, to one in which the identifying content and the useful content are close to the same thing.

6 Dual Use and the Question of the Operator

The strongest objection to the proposal is that a publication projection implemented in software restricts what a system says and does not restrict what the apparatus can observe or store. A camera that emits hazard events is, physically, a camera. With a change of firmware it can read plates, track vehicles, or retain footage, and nothing in the public-domain status of its hazard feed prevents that change. The legitimacy of the arrangement therefore depends on matters that the formal model leaves outside itself: on what the hardware is capable of capturing, on who controls the firmware, on whether the running configuration can be audited, on what retention policies bind the operator, and on what institutional authority can compel the operator to change any of these.

The four properties distinguished earlier bear directly on this. Output minimization

can be demonstrated from outside, since anyone can read the feed. Non-retention and verifiable enforcement cannot be demonstrated from the feed alone, because the feed is silent about what the device does with what it does not emit. Remote attestation of firmware can establish that a particular signed configuration was running at a particular time, which is evidence of a state and not a commitment about the future. Hardware that physically lacks the resolution or the storage needed to capture identifying detail secures non-acquisition, which offers a stronger assurance than any of the operational properties, because the capability is absent and not merely disabled, but such hardware is a design choice that must be made and maintained, and it can be replaced. Independent audit, public disclosure of retention rules, and legal limits on repurposing each contribute something, and none of them converts a promise about behavior into a property of the device.

A further complication concerns the operator. The institution that deploys the sensor has the technical means to turn it into a surveillance instrument, and the credibility of its stated restriction rests on its own willingness to be bound. Where the operator is also the party that would benefit from surveillance, an output restriction that the operator alone can lift is weaker than it appears. Where the operator is subject to external constraint, with the capacity of outsiders to inspect, contest, and sanction, the restriction acquires force that no technical feature supplies. The mathematics of the earlier sections is indifferent between these cases. The difference between them is institutional, and it is where the argument either succeeds or fails.

The essay has argued that the public domain should be defined by what a system is authorized to release, that this authorization is best understood as a projection with a stated purpose, that the formal statement of the projection settles nothing about who chooses its parameters, and that the strength of the resulting privacy claim varies widely across cases. What it leaves unresolved is the question on which the arrangement ultimately depends: whether an institution can legitimately claim to operate a publication-only instrument when it retains the technical capability to turn that instrument into a surveillance device.

References

- [1] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, 2006.
- [2] M. Gymrek, A. L. McGuire, D. Golan, E. Halperin, and Y. Erlich. Identifying personal genomes by surname inference. *Science*, 339(6117):321–324, 2013.
- [3] N. Homer et al. Resolving individuals contributing trace amounts of DNA to highly complex mixtures using high-density SNP genotyping microarrays. *PLoS Genetics*,

4(8):e1000167, 2008.

- [4] A. Makhdoumi, S. Salamatian, N. Fawaz, and M. Médard. From the information bottleneck to the privacy funnel. In *IEEE Information Theory Workshop*, 2014.
- [5] A. Narayanan and V. Shmatikov. Robust de-anonymization of large sparse datasets. In *IEEE Symposium on Security and Privacy*, 2008.
- [6] L. Sweeney. Simple demographics often identify people uniquely. Carnegie Mellon University, Data Privacy Working Paper 3, 2000.
- [7] N. Tishby, F. C. Pereira, and W. Bialek. The information bottleneck method. In *Proceedings of the 37th Allerton Conference on Communication, Control, and Computing*, 1999.