Opacity as Constitutional Capacity: State Legibility and Privacy Technology

Eric Alston - University of Wyoming
Federica Carugati - King's College London

Abstract: We develop a constitutional theory of the relationship between three distinct features of the state-citizen relationship: legibility, defined as what the state can reliably observe and act upon; privacy, defined as the protected sphere within which individuals act outside full state visibility; and opacity, defined as the level of effective non-visibility that obtains in a given society at a given time. We characterize this opacity level as a function of the institutional, technological, and individual forces that determine where on the privacy/legibility continuum a given polity is located. We argue that legibility and privacy are mutually constitutive components of democratic capacity, and that the opacity that emerges from their interaction is peculiarly sensitive to changes in information technology. Earlier step-changes in information capacity moved the opacity balance toward legibility, but at a rate that ordinary institutional correction could absorb. The digital era is different in three respects: legibility is already high, technological shocks to government capacity arrive in compressed succession, and the corrections that would maintain a democratically-representative opacity balance are lagged, partial, and increasingly bypassed by commercial data acquisition and surveillance by executive enforcement authorities. The resulting drift is a constitutional dilemma, in which a democratic state's observed behavior can diverge from the privacy commitments it has formally assumed. Privacy technology has become a corrective on the privacy side of the continuum, because it operates at the speed of adoption rather than the speed of public governance. The opacity balance is therefore a constitutional question because it depends on tradeoffs among competing legitimate governance preferences that only democratic contestation can weigh, and a polity that lets its opacity drift by administrative default has not struck a balance but instead abdicated the choice.

Introduction:

We develop a theory of the relationship between three features of the state–citizen relationship: legibility, what the state can reliably observe and act upon; privacy, the protected sphere within which individuals act outside full state visibility; and opacity, the level of effective non-visibility that actually obtains in a society at a given moment. Privacy and legibility stand at opposite ends of a continuum, and a society's opacity is the point it occupies on that continuum. This opacity is the observed result of the state's need for legibility, the pace of technological change, and citizens' demand for privacy, including their ability to adopt technology that preserves their privacy. Opacity, rather than legibility or privacy considered alone, better captures the observed condition citizens actually live under and the one against which a state's commitments are tested.

We argue that privacy and legibility are mutually constitutive components of democratic capacity rather than opposing values to be traded off. Legibility is a precondition to the state's capacity to tax, to enforce, to coordinate, and to protect. Privacy is an input to the legitimacy on which state capacity ultimately depends. A state that expands its visibility without limit does not thereby maximize its capacity; past a threshold it erodes the consent, cooperation, and trust that improve capacity. Legibility and privacy are thus bound together, and the balance between them that determines a society's opacity is peculiarly sensitive to technological change.

That sensitivity is the source of the problem we identify. Historically, states have strengthened administrative and fiscal capacity through projects of legibility such as censuses, standardized identities, and cadastral surveys (D'Arcy and Nistotskaya 2017; Lee and Zhang 2017; Brambor et al. 2020; Rogowski et al. 2021; Christensen and Garfias 2021; Bowles 2024). Legibility increases implementation efficacy, but excessive legibility may weaken legitimacy, as well as undermine adaptability, provoke resistance, or generate fragile institutions (Scott 1998). Constitutional democracies answered the resulting risks by institutionalizing spheres of privacy, the informational counterparts of the property-rights and rule-of-law commitments that accompanied the modern democratic state (North and Weingast 1989; Acemoglu and Robinson 2012). Yet these privacy commitments were built for an environment of physical search, paper records, and analog wiretaps. Digital technology has utterly altered that environment, and has done so faster than the commitments could adapt. The result is a ratchet: each cycle of innovation raises legibility, while the corrections that would restore privacy arrive late, reach only forward, and leave in place both the data already gathered and the enforcement actions already taken on it (Collingridge 1980; Downes 2009).

Despite the legitimate justifications for a state to "see" its citizens, unrestrained legibility can be self-defeating. When people know they may be extensively monitored, they self-censor and withdraw from association and participation (Penney 2016; Marthews and Tucker 2017); comprehensive surveillance corrodes the public trust on which voluntary compliance and cooperation depend (Cingolani 2023; cf. Levi 1988); and as the state leans on opaque automated

systems, it compromises its own accountability (Bloch-Wehba 2021). Legibility pursued past the point the polity would authorize can thus erode the very capacity it was meant to extend.

This makes any such persistent drift toward legibility a constitutional dilemma. Because of this, where the opacity balance should sit is an irreducibly political question: it turns on tradeoffs among competing and legitimate citizen preferences and no expert answer can install the right mix in place of a representative one. A balance set instead by the accumulated momentum of executive, enforcement, and administrative decisions, each sensible on its own terms, is not a balance the polity chose; it is a de facto answer to a question that was never put. To call the balance constitutional is to insist that it belongs to the polity as a whole, settled through plenary democratic contestation rather than fixed by default through administrative accretion.

This insistence becomes more urgent as a given state proves less able to maintain the opacity balance through ordinary channels. Where public institutions cannot keep pace, the supply of privacy comes to rest substantially on citizens themselves and on the producers of privacy-enhancing technology: strong encryption, peer-to-peer infrastructure, pseudonymity, shielded payment networks. That these tools are increasingly functionally essential is exactly why their governance cannot be left to default. If the state cannot reliably regulate its own legibility, the question of what privacy technology a polity should permit, restrict, or require is one democratic contestation must answer.

It is also why a state's stance toward privacy technology is so revealing. This is the domain in which the incentives of the actors who wield the state's enforcement and administrative authority cut most reliably against privacy: states have a standing reason to see more of what its citizens do, and technologies that place activity beyond its view press directly against that reach. How a state's institutions behave here therefore reveals two things at once — its substantive position on the opacity balance, and the seriousness with which its institutions honor the commitments they profess when those commitments are most costly to keep. A state's conduct toward privacy technology is, in this sense, a bellwether: of where it sits on the continuum, and of the fidelity of its constitutional commitments more generally. In the current period, conversations about privacy technology belong at the center of democratic debate because they reveal this balance, and the open question for any constitutional democracy is which responses to the expansion of legibility it should deliberately institutionalize, rather than allow to accrete unchosen.

I. Privacy, Legibility and Opacity

Our argument turns on three concepts: privacy, legibility and opacity. Privacy refers to the protected sphere within which individuals and groups act outside full state visibility. Legibility refers to what the state can reliably observe, measure, and act upon. Opacity refers to the level of effective non-visibility that obtains for individuals and organizations in a given society at a given time. If we take privacy and legibility as opposite ends of a continuum, a society’s level of

opacity is the point on that continuum the society occupies at a given time. This understanding follows from the institutional economics tradition that frames privacy as an equilibrium outcome of strategic interaction over information (Posner 1978; Stigler 1980; Hirshleifer 1980; Murphy 1996; Hermalin and Katz 2006; Acquisti, Taylor, and Wagman 2016). Once privacy is so framed, the observed level of opacity at any given time becomes an empirical question (Acquisti, Brandimarte, and Loewenstein 2015; Penney 2016), but also enables us to cast privacy and legibility as a continuum in light of these tradeoffs. Any point on this continuum is the product of the actions of individuals and states, given certain information technology affordances.

First, we should clarify that we use the term state as ordinary shorthand, but this singular noun should not be mistaken for a single actor with a single will. A modern state is the emergent output of legislators, courts, agencies, investigators, prosecutors, and administrators, each pursuing competing, overlapping, and sometimes conflicting aims within a set of rules, and what the state "does" at any moment is the result of that interaction rather than the choice of one mind. This matters because the dynamics of privacy and legibility that we describe are not reducible to the demand of a unified sovereign and instead result from predictable pressures that bear on particular classes of actors within the “state.” What differs across states is precisely how well the specific institutional arrangement, taken as a whole, channels and checks those pressures. When we speak of a state's "stance" or "conduct," we mean an observable position that emerges from this interaction among numerous actors, not a preference the state rationally holds.

Understanding the relationship between legibility and privacy as a tradeoff (Posner 1978; Stigler 1980) means that holding the technological stock fixed, more of one typically means less of the other. Holding the technological stock fixed, more legibility of a given kind means less privacy of that kind: more observation of citizens' communications is less communicative privacy, more observation of their transactions is less financial privacy, more observation of their movements is less locational privacy. This is a structural feature of every state as opposed to a contingent feature of any particular regime. Legibility and privacy, then, cannot be jointly maximized. This is the structural feature that gives the institutional privacy dilemma its bite. However, for states that have never committed to privacy affordances for their citizens, the dilemma is less present, even if their citizens might prefer a regime that offered such affordances, however imperfectly.

Because of the structural tension we have developed here, it is more productive to focus not on legibility or privacy separately but on the level of opacity that obtains in a given society at a given moment. Opacity is the operative continuum: the actual, observed level of effective non-visibility citizens enjoy in their affairs, taking into account constitutional commitments, statutory regimes, judicial interpretation, administrative practice, the available technological stock, and the privacy-related choices of citizens themselves. A society at high opacity is one in which the legibility/privacy balance tilts toward privacy; a society at low opacity is one in which it tilts toward legibility.

The same level of opacity can be reached through different combinations of these inputs: a society with weak privacy law but strong privacy technology may sit at the same opacity level as a society with stronger privacy law but weaker technological adoption. What matters for citizens, and for the constitutional analysis we develop, is the resultant level rather than the inputs taken in isolation. As we define opacity, it captures the aggregate condition that bears on citizen welfare, democratic legitimacy, and the credibility of state commitments, without prejudging which mix of institutional and technological forces produces it, and what the "right" level for a given society is. It is also why opacity, rather than legibility or privacy considered separately, is what the rest of this argument will track. The institutional privacy dilemma is, properly understood, a dilemma about the level of opacity that public institutions can credibly sustain in the face of competing pressures and shifting technological capabilities.

Legibility, in the sense we adopt here, is a term of art borrowed from the literature on the informational foundations of state capacity (Foucault 1975; Scott 1998). The concept's roots lie in the distinction between despotic and infrastructural power: a state's infrastructural capacity to penetrate society and implement decisions across its territory is realized by the administrative and informational apparatus that renders populations visible and actionable (Mann 1984; Tilly 1990; Soifer 2008). Scott's account of how modern states render populations and territories legible through innovations such as cadastral surveys, standardized identities, and administrative classifications established the basic point that the state's ability to act effectively is a function of what it can see (Scott 1998). More recent work in historical political economy and comparative public administration has formalized this insight (Besley and Persson 2009; D'Arcy and Nistotskaya 2017; Lee and Zhang 2017; Brambor et al. 2020; Rogowski et al. 2021; Christensen and Garfias 2021; Bowles 2024; Cingolani 2023). The cumulative finding is that legibility is not a downstream effect of state capacity but a constituent component of it. A state without reliable information about its citizens cannot administer them, tax them, protect them, or provide them with public goods.

Privacy, by contrast, is the condition of being unobservable, untracked, or unmeasured by the state in some defined respect. The concept is older than legibility, and we do not attempt a full accounting of its philosophical or doctrinal dimensions (but see Warren and Brandeis 1890; Westin 1967; Gavison 1980; Solove 2006, 2008; Nissenbaum 2010).[1] Our approach to privacy combines Westin's informational-control framing as a baseline (privacy is about who knows what) but also adopts Solove and Nissenbaum's resistance to a single unifying essence (privacy is pluralistically constituted, and opacity is the operative result rather than a single condition). In the last several decades, privacy has come into renewed analytical focus because the dramatic growth of information technology capabilities has brought new challenges to the fore. The digital

[1] Warren and Brandeis (1890) treated privacy as a tort-based 'right to be let alone'; Westin (1967) recast it as informational control; Gavison (1980) disaggregated it into secrecy, anonymity, and solitude as distinct conditions; Solove (2006, 2008) developed a pluralist taxonomy of sixteen privacy-implicating activities resistant to single-essence definitions; and Nissenbaum (2010) reframed privacy violations as breaches of context-specific informational norms ('contextual integrity').

era's combination of cheap sensing, near-zero-cost replication, automated inference, pervasive cross-dataset linkage, and rapid commercial procurement has produced a step change in what the state can practically observe about its citizens (Solove 2008; Cohen 2013; Richards 2015; Zuboff 2019). The change has generated a deepening asymmetry between the relatively stable institutional commitments to privacy that states inherited from the analog era and the rapidly expanding informational reach of the contemporary state.

How did the balance between legibility and privacy evolve during past ICT revolutions? Little empirical work has been carried out to answer this question. However, the topic has received a great deal of attention in history, archeology, anthropology, and sociology.

Take, for example, the invention of writing: long standing archeological and anthropological accounts and limited empirical work converge in suggesting that writing dramatically extended the state's ability to see and record people's movements and activities (Childe 1936, 1950; Goody 1986; Nissen 1993; Yoffee 2005; Hudson 2000; Wang 2014; Stasavage 2021; Benati, Carugati and Leon-Ablan 2026). Whether technological adoption was welcome or met with resistance, however, is harder to infer from the archeological and historical record. The first evidence from Mesopotamia appears almost a millennium after the invention of writing in Uruk around 3300 BCE, when it took the form of exit: people relocated to more remote areas to avoid taxation, corvee work, and conscription (Richardson 2017). Writing was a technology that spread "from" the state, as scribes were first trained and employed by states, especially religious elites governing temple-based polities (Liverani 2006; Benati, Carugati and Leon-Ablan 2026). Things were different, however, with the next ICT revolution: the printing press.

Early printing developed first in China in the 7th century from Buddhist monks seeking to spread their ideology (Miller 1983; Cheng et al. 2023). The technology was only later adopted by the state to set the imperial examination that would grant entry to the nascent Chinese bureaucracy. Religious competition may have played a role in Europe too, given the role of the technology in the spread of Protestantism in the 16th century. But whether European states experienced a similar indirect capacity-enhancing effect as their Chinese counterpart is an understudied topic. Empirical studies especially in economic history have concentrated on the impact of the printing press on economic growth and on the spread of the protestant reformation, but not on state capacity (Dittmar, 2011; Rubin 2014, 2017; Boerner et al. 2021; Becker et al. 2023). An inference, however, can be drawn indirectly: much of the later work on information capacity we discussed above relies on the existence of printing technology, as censuses, registries, and public agencies rely on reproducible documents, standardized formats, and dissemination technologies (see also Gadd and Gillespie 2006; Mukerji 2011).

To better understand the effects of mass printing on states we rely on the work of historians and sociologists who have documented these effects in greater detail. These accounts suggest that states used the technology - much as they did in later ICT revolutions - to control mass

communication through both censorship and propaganda. Many forms of resistance developed against the state's reach on the press, from pamphleteering to underground printing networks generating new forms of mass protest and dissent, not infrequently leading to violence.[2]

These strategies of resistance to increases in state legibility will not surprise a social science audience familiar with the work of James Scott (1985; 1998; 2009). But, we argue, the pace of technological change in the past gave state and non-state actors a chance to match (over)reach and contestation. Indeed, the most flagrant abuses led to dramatic outcomes for the contestants (Scott 1998). But in many cases, contestation over legibility projects enhanced state capacity and, at times, freedom of information. An example of this is the case of licensing laws in 17th century England (Raymond 1996; Black 1987).

In this paper, we suggest that this dynamic, which we describe later as the "static" version of our theory, no longer applies. The asymmetry of the present technological era makes understanding the balance between legibility and privacy newly urgent.

What distinguishes the modern era from each of these prior step-changes in information technology exceeds the magnitude of any individual shock because this era combines three features: the already-high baseline legibility against which new shocks register, the compressed interval between shocks, and the increasing scale and granularity of the capabilities each shock unlocks. Even though modern constitutional democracies were built on the recognition or spheres of privacy, the privacy-protective doctrines and statutes built for an environment of physical search, paper records, and analog wiretaps were not designed for, and do not easily map onto, an environment in which most of a citizen's daily activity generates granular and wide-ranging personal data available to commercial intermediaries from which the state can purchase or subpoena it.

II. Privacy, Legibility, and the Pace of Technological Change

Following the distinction between legibility, privacy, and opacity, we turn to how the balance between privacy and legibility responds to change in information technology. Our argument is that a constitutional problem emerges only when three conditions hold. First, the static institutional privacy dilemma gives a wide range of actors within the state standing incentives to expand legibility, so that new capabilities are adopted before they have been constitutionally legitimated, generating a continual need for correction. Second, those legislative and judicial corrections, though they err in both directions, produce errors that are not equally durable: an error toward legibility tends to persist while an error toward privacy is readily unwound, so that

[2] There is some limited literature on the impact of the telegraph on state capacity, including on infrastructure in Chile (Martland 2014) and on coercion in the US (Kielbowicz 1994; Parinandi 2023), but the findings are not systematized. For brevity, and because our primary goal is not to document the development of ICTs in history, we skip this revolution whose effects and countereffects, we suspect, would be similar to those we discuss for the printing press.

even an unbiased corrective process drifts in one direction. Third, as the rate of technological change rises, these persistent errors begin to compound, each accumulating on the residue of the last faster than correction can respond. These three conditions correspond to the level of opacity at a moment, the rate at which that level changes, and the rate at which that rate itself changes. The constitutional problem is a property of the third: it arises not from the existence of a ratchet effect but from its acceleration past the point where democratic correction can keep pace. We take the three in turn.

A. The Static Institutional Privacy Dilemma

Treated statically, the institutional privacy dilemma is intuitive. The state is responsible for defining and enforcing privacy protections, yet the same state need for legibility requires informational access to perform core governance functions: criminal investigation, taxation, border control, regulatory enforcement, national security. The public actor expected to protect a right to privacy is also uniquely incentivized to infringe it (Alston and Cossar 2026). From this structural tension, the static analysis yields an observed level of opacity in a given democratic order at a moment in time, in which constitutional protection, statutory authorization, judicial review, and administrative procedure jointly constrain the scope of state visibility while preserving the informational foundations of governance.

It is important to recognize that this balance between privacy and legibility is both a contest between citizen preference and state interest, as well as among competing citizen preferences themselves (Etzioni 1999; Allen 2011). Many of the same citizens who value privacy also benefit from the state capacity that legibility enables: enforcement that protects them when victimized, taxation that funds public goods , and administrative coordination that delivers services. Sufficient legibility is to some extent required for a state to provide what citizens demand when they demand effective government. Many consequential legibility expansions of the last half century have been adopted not over citizen objection but in response to citizen demand: for example, AMBER Alerts, sex offender registries (Logan 2009, 2011), modernized tax administration that reduces compliance burdens for honest filers, and the routine use of automated license plate readers to recover stolen vehicles and locate missing persons. A state's level of opacity on the privacy/legibility continuum is correspondingly contested in a deeper sense than our earlier framing suggests: contested between competing legitimate preferences individual citizens may hold, and across heterogeneous populations (Solove 2021). Where a state has explicitly committed itself, through constitutional text, jurisprudential interpretation, or treaty accession, to honor citizen privacy, the work of legislative deliberation, judicial interpretation, and administrative implementation is to give content to that commitment over time; this political contestation is necessarily costly, time-consuming, and contentious, and yet under the right conditions produces outcomes recognizable as representative, at least in a relative sense compared to regimes that do not make commitments to representative legitimation.

A static account of the privacy dilemma is not wrong on the face of it, because the level of opacity citizens enjoy in their affairs directly informs citizens' perceived right to privacy and the extent to which the state can preserve it. This static account captures the level of opacity at a moment in time, but it is silent on how that level moves. It is only once we consider the rate at which opacity changes, and then the rate at which that rate itself changes. that the constitutional dimension of the dilemma comes into view. We take these two intertwined yet analytically distinguishable dynamics in turn.

B. Legibility Increases Precede Imperfect Corrections

Consider first the rate at which opacity changes, setting aside for the moment the question of whether that rate is itself rising. Even at a constant pace of technological change, the relationship between legibility shocks and privacy corrections is asymmetric: legibility-enhancing capabilities arrive as discrete shocks, adopted quickly and at low marginal cost, while the corrections that would restore a representative opacity balance arrive only through slow, ex post institutional channels such as statutory reform, doctrinal evolution, evidentiary challenge, oversight review, and administrative policy. But the problem is not only that correction is slow. It is that correction, when it comes, is imperfect in three distinct ways, and the imperfections compound rather than cancel.

It is worth being explicit that the dynamics that we describe here as corrective are necessarily error-prone in both directions. First, privacy-restoring responses to a legibility shock have observably succeeded: Carpenter, Riley, and the USA FREEDOM Act each restored some measure of privacy that prior practice had eroded. Second, and relatedly, corrections can err in either direction: a court or legislature can constrain the state more than the representative balance warrants, not only less. We do not claim that corrections are systematically too weak, nor that any given correction cuts toward legibility. Our narrower claim is accordingly stronger: corrections are a genuine two-directional instrument, but the errors these corrections entail are not equally durable. An over-correction toward privacy is more readily unwound, whether through renewed collection, later revision, a subsequent administration or exogenous technological change. Yet an under-correction toward legibility tends to persist, because the data already gathered survives the rule that should have reached it and flows beyond the boundary that should have held it. Even an unbiased corrector, making as many errors toward privacy as toward legibility, therefore yields a system that drifts in one direction, because the system keeps the legibility errors and discards the privacy ones. It is this asymmetry of persistence in corrective errors that the three issues we now detail indicate.

The first corrective deficiency is temporal. During the interval before a correction arrives, the uncorrected balance is not idle. Administrative determinations are made, enforcement actions are taken, prosecutions are brought, and people are acted upon under a level of legibility that no representative process authorized. These are not costs deferred until the correction restores the

balance; they are realized as they occur and, for the most part, irreversible. A warrant requirement announced after a decade of warrantless collection does not reopen the cases the collection produced, and a doctrinal limit on a surveillance practice does not undo the determinations already made on its fruits. Each unit of delay is thus a stock of state action taken at an incorrect point on the continuum, and that action stands even after the rule that permitted it is changed. The lag should be understood as time during which the wrong opacity balance produces consequences that remain.

The second corrective deficiency is one of reach: a privacy correction operates on the rule going forward and rarely reaches the stock already accumulated. A judicial or statutory limit on collection constrains what the state may gather next, but the data already gathered, copied, modeled, and networked into interagency systems is seldom expunged by the same act that limits its future growth. Carpenter constrained prospective cell-site collection without unwinding the records already held; the USA FREEDOM Act curtailed future bulk collection while leaving the assembled metadata corpus in place. The rule turns forward; the stock reaches back. This asymmetry is reinforced by a second, between the two kinds of rule themselves. A rule that expands legibility rides existing incentives and accrues operational dependence across budgets and vendors, which makes it durable and resistant to repeal; a rule that constrains legibility runs against the interest of the actors it binds and against the inertia of an infrastructure already built, so that even a well-designed protection must survive continuing pressure from the very incentive it exists to check. Durable privacy protections are therefore possible and matter. Nonetheless, they are harder to win and harder to hold than the legibility expansions to which they respond, and even when they hold they leave the prior stock untouched.

The third corrective deficiency is one of containment: private information, once held, does not stay within the authorization that gathered it. Because the state is not a single actor but many actors with overlapping access to common data, a rule that constrains one actor's use of information does not reliably constrain every actor who can reach it. Information lawfully collected by one set of actors for one authorized purpose flows to another set pursuing a different one, across the very boundaries a rule was meant to hold. The practice of parallel reconstruction is the covert version: investigators build an alternative evidentiary trail to hide the surveillance origin of a lead, suppressing the records on which doctrinal correction depends (Babazadeh 2018). This presents a deeper irregularity, since the information “laundered” was often collected under one authority, frequently foreign-intelligence collection carrying predicates distinct from those of ordinary criminal investigation, and the information is then put to a use that authority was not authorized for. Recent litigation alleging that state police in Oregon allowed federal immigration authorities to query state databases for years, in tension with the state's sanctuary law, is the overt version of the same failure: a legislatively enacted rule constraining the use of an existing stock eroded not by repeal but by the ordinary interoperability of the systems that hold the data.[^n] The two are one phenomenon seen from opposite sides — in one the breach is rendered invisible so it is never adjudicated, in the other it surfaces and is litigated — and in

both, the leakage is not the decision of a unitary state defying its own commitment but the emergent product of many actors with access to a common stock, which is precisely why a rule directed at one of them does not bind the rest.

These asymmetries are not a sign of institutional failure. The channels through which corrections occur are slow because they are ideally careful, weighing competing legitimate preferences across heterogeneous populations through institutions whose own legitimacy depends on procedural regularity. But the same features that make correction arguably more reliable when it eventually arrives also make it lag, often by years, and the result is a structural mismatch: legibility expands in discrete, rapid increments, while the privacy-restoring response arrives late, partial in reach, and porous at its boundaries.

The post-9/11 surveillance trajectory illustrates this asymmetry (Donohue 2016). Section 215 of the USA PATRIOT Act was enacted in October 2001 and reauthorized repeatedly through the next decade and a half. Bulk telephony metadata collection under that authority was operational by 2006 and would not be publicly known until Edward Snowden's June 2013 disclosures, by which point the FISA Court had reauthorized the program more than thirty times under fourteen different judges. The legislative correction in the USA FREEDOM Act did not arrive until June 2015, nearly fourteen years after the underlying authority and nine years after the program itself began. Even then, the correction was partial: it replaced bulk collection with a narrower call detail records regime that, in 2018 alone, generated more than 434 million records under just one of its two Section 215 provisions, and it left in place much of the broader architecture of FISA Section 702 and the Foreign Intelligence Surveillance Court, which retains a structurally weak adversarial process even after USA FREEDOM Act amicus reforms (Goitein and Patel 2015; Vladeck 2015).

The doctrinal record displays the same asymmetry. The third-party doctrine announced in United States v. Miller (1976) and Smith v. Maryland (1979) survived essentially intact for four decades; its first significant constitutional retrenchment came in Carpenter v. United States (2018), and even that decision limited only historical cell-site location information beyond a seven-day window, leaving real-time tracking, shorter-duration requests, and the underlying logic of the doctrine for non-location records substantially unchanged. Riley v. California (2014) imposed a warrant requirement on cell-phone searches incident to arrest only after roughly a decade in which such searches had been treated as routine. The Department of Justice for years operated cell-site simulators ("Stingrays") that recorded adjacent cell phone activity before issuing its first internal policy in 2015 requiring warrants for their use, and even then the policy bound only federal agencies. The geofence warrant, a search technique that grew from a handful of requests to Google in 2016 to more than 11,000 in 2020, reached the Supreme Court only in Chatrie v. United States, argued in April 2026, after Google itself had already restructured its data storage to render the technique substantially obsolete. In this last case, technological change obviated the very practice the Court had agreed to review.

This last point illustrates the rebound problem in its purest form, but the deeper lesson of the record is cumulative: each of the three failures is visible in it. The years before each correction were years of realized enforcement under an unauthorized balance; the corrections that came reached forward to the rule and left the accumulated stocks in place; and those stocks, once held, propagated across the agencies and authorities able to reach them. Opacity drops sharply when legibility tools are adopted, and tends to recover only partially, late, and leakily when corrections eventually arrive. Yet this leaves aside the crucial question of the rate at which these shocks arrive within a particular democratic order, and the corresponding question of where on the privacy/legibility continuum that order was situated when the shocks began. To answer either requires a step back from the contemporary case, because legibility drift in the digital era does not operate against a neutral baseline. It operates against a historical trajectory in which earlier step-changes in information technology moved the opacity balance toward legibility at a cadence slow enough that contestation and institutional adjustment could, however imperfectly, keep pace. In contrast, the modern condition of high baseline legibility and rapidly compounding technological change is itself a recent departure, whose implications we next consider.

C. Accelerating Technological Change and the Drift Toward Legibility

We come now to the rate at which the rate itself changes. The ratchet described in the previous subsection would be tolerable, and largely self-correcting, if the pace of legibility-enhancing innovation were constant: institutions could catch up between waves, and whatever balance contestation reached had a reasonable chance of holding. The distinctive feature of the digital era is that the pace is not constant but accelerating. Cheap sensing, ubiquitous data collection, near-zero-cost replication, automated inference, pervasive cross-dataset linkage, and rapid procurement through private vendor ecosystems compress the interval between legibility shocks.

The next information capacity shock may arrive even before the institutional response to the last has been fully countenanced.The contemporary deployment of surveillance infrastructure illustrates the compressed cadence of an increasing rate of innovation in information technology. Flock Safety, founded in 2017, has placed roughly 90,000 automated license plate readers across approximately 6,000 municipalities in nearly every U.S. state, generating something on the order of twenty billion plate reads per month and pooling that data into a nationally searchable network. Flock represents a near-total fielded infrastructure assembled over less than a decade, while no comprehensive federal statutory framework yet governs its use and the first state-level comprehensive ALPR statute, Oregon's SB 1516, took effect only in March 2026.

Similar dynamics characterize commercial data acquisition. U.S. Immigration and Customs Enforcement has assembled a deportation infrastructure built largely from purchased records that includes coverage of driver location data in cities home to roughly 75 percent of the U.S. population, utility records that locate three of every four adults, and access to driver license data for a comparable share (Wang et al. 2022). Setting aside the specific public administrative uses

of state and local data, much of this data is acquired through commercial intermediaries in a manner that bypasses Fourth Amendment constraints altogether by virtue of having been collected by private parties in the first instance (Slobogin 2008). Meanwhile, vendors like Clearview AI assembled facial-recognition databases of tens of billions of scraped images before any sustained doctrinal or legislative engagement with the practice occurred (Garvie, Bedoya and Frankle 2016; Selinger and Hartzog 2019). Against each of these positive information capacity shocks, legislative correction has been jurisdictional, partial, and slow; doctrinal correction has been case-bound and often years behind operational deployment.

This is the threshold at which the analysis shifts from first order to second. A ratchet operating at a constant rate produces a predictable, and therefore correctable, decline in opacity; institutions know what is coming and can, in principle, keep pace. But once the rate of legibility expansion is itself rising (i.e., once the shocks arrive closer together and each leaves a residue from which the next compounds) the decline in opacity accelerates, and correction that was merely lagging becomes structurally unable to catch up. It is here, at the second moment, that drift becomes not merely possible but structural: where a constant-rate ratchet might still be outrun by a sufficiently weak corrector, an accelerating one outruns any correction sooner or later. What our stylized model[3] identifies is a structural tendency: as the rate of information capacity shocks increases, the probability that a given system drifts beyond its representative opacity balance (wherever that balance happens to lie) increases, all else equal. The static institutional privacy dilemma is not displaced; its momentum is defined by a dynamic asymmetry in the direction of increasing legibility. As argued above, faster cycling does not wash the drift out: because legibility errors persist and privacy errors are unwound, running more error through the system in less time compounds rather than cancels the imbalance.

This dynamic is amplified by a feedback effect. Some legibility-enhancing innovations also reduce the contestability of state information practices. Proprietary systems, complex automated pipelines, secrecy norms around vendor relationships, and the sheer scale of commercially available data acquisition lower the visibility on which evidentiary, doctrinal, and oversight correction depends. Standing doctrine itself has narrowed the range of plaintiffs able to challenge surveillance programs whose existence they cannot prove with concrete particularity, as the Court's 2013 decision in *Clapper v. Amnesty International* illustrates (Vladeck 2014). Trade-secret protection over proprietary risk-scoring and pattern-detection algorithms, of which *State v. Loomis* (Wisc. 2016) and the broader litigation surrounding the COMPAS recidivism instrument are well-known examples, has insulated automated state decisions from meaningful evidentiary challenge. The practice of reconstructing alternative evidentiary trails to launder the

[3] Let O(t) denote the level of opacity at time t. The static dilemma concerns the level O(t); the ratchet effect concerns its first derivative $dO/dt < 0$, a directional decline in opacity as legibility shocks outpace corrections; and the acceleration concerns the second derivative $d^2O/dt^2 < 0$, a decline whose rate is itself increasing. The constitutional problem we identify is a property of the second derivative: drift in the strong sense arises when the magnitude of d^2O/dt^2 exceeds the throughput of institutional correction, so that the gap between observed and authorized opacity widens faster than deliberation can close it.

surveillance origins of a prosecution's lead directly suppresses the case records on which doctrinal correction depends (Babazadeh 2018). As contestability falls, correction itself becomes weaker, and the political work of locating the "right" opacity balance becomes harder to perform. As correction weakens, the next legibility shock encounters a more permissive institutional environment than the one before. The ratchet effect is therefore not only a function of relative speeds; it is also a function of an endogenous decline in the system's own capacity to push back, and, more fundamentally, in its capacity to know whether it should.

Put directly, the drift we describe can become a representativeness failure past a certain level of opacity. For political contestation over the opacity balance to produce outcomes that can plausibly be called representative, one of two conditions must hold (Habermas 1996; Mansbridge 2003). Either public institutions must not exhibit a systematic tendency to drift in one direction independent of citizen preference, or, where it does drift systematically, the direction of drift must correspond to the privacy/legibility preferences of a sufficient majority of the polity. The drift here violates the first condition by definition: the asymmetry between rapidly-arriving legibility shocks and lagged, partial corrections produces a one-directional tilt toward legibility regardless of the politically-contested balance that would otherwise have been struck. Whether it also violates the second condition is an empirical question about citizen preference, but the structural argument does not depend on resolving it. What matters is that the burden of demonstrating majority ratification of the drift falls on those defending the resulting balance, and that burden becomes harder to meet as the gap between the politically contested balance and the actual level of opacity widens. Where neither condition holds, the level of opacity that obtains in society at any moment is not a representative outcome of political contestation; it is the residue of a process in which one side of the contest has structural advantages the other does not. The republican tradition adds a further reason this matters: even an unused expansion of legibility constitutes a form of domination, since freedom in this conception requires the absence of arbitrary capacity for interference instead of just the absence of its actual exercise (Pettit 1997; Roberts 2015). Surveillance produces an inequality of power between watcher and watched that is itself a representativeness asymmetry (Richards 2013). This is why the legibility drift's harm is constitutional rather than regulatory. It corrodes the conditions under which political contestation can be said to authorize the balance it produces.

III. The Structural Role of Privacy Technology

We have argued that public institutional channels cannot reliably maintain a representative opacity balance against accelerating legibility shocks past a certain point. This raises the question of how a representative balance can be approached at all under such conditions, and we argue that this depends on the technological stock available to citizens (end-to-end encryption, virtual private networks, and shielded payment networks), which partly determines available opacity at

the speed of adoption rather than legislation, and whose development has to date been largely independent of the apparatus of political contestation.

This dependence is not uniform across states; it is a function of how reliably a state's own commitments hold. Where institutions correct faithfully and promptly, privacy technology is an affordance citizens may or may not adopt. But where a state's privacy commitments are structurally weak, or have proven systematically defective in practice, the technological stock becomes a principal guarantor of whatever privacy citizens retain. The more defective the institutional supply, the more load-bearing the private one becomes. This normative political question of the "right" opacity balance is separable from where the supply of privacy actually comes from when the state's own provision fails. This makes the legal status of privacy technology, such as whether citizens may adopt it, and on what terms, a question of constitutional rather than merely regulatory import.

Our structural argument applies across constitutional democracies, but examining how a state's stance toward privacy technology raises constitutional questions requires granular institutional detail, and we focus on the United States for three reasons: it has explicitly constitutionalized privacy commitments, principally through Fourth Amendment doctrine; its broadly permissive baseline makes it an illustrative case, since episodes of suppression there read as revealed preference of state actors rather than the expression of a more generally restrictive regime; and its record over the last decade varies across distinct technologies and successive administrations of both parties, letting us observe whether the direction of pressure is observably invariant to the technology and the government applying it.

A. Privacy tech is central to the opacity balance in the digital era.

The technological stock that bears on the opacity balance is largely exogenous to the process of political contestation. Neither citizens nor legislators directly produce or suppress the underlying capabilities. Encryption protocols, shielded payment networks, distributed-ledger architectures, and zero-knowledge constructions are developed within open-source communities, academic cryptography, and competitive commercial environments whose outputs diffuse globally once published. The same is true on the legibility side: the sensors, models, and data infrastructures that increase drift are produced by private vendors, foreign actors, and decentralized research communities whose work is not contingent on any particular jurisdiction's authorization. The general phenomenon, that technical architectures themselves operate as a form of regulation orthogonal to and often beyond the reach of statutory command, is well established in the law and technology literature (Lessig 2006; Reidenberg 1998).

Where public institutional correction is too slow relative to the rate of legibility shocks, the burden of preserving citizens' capacity to maintain spheres of privacy falls on the technological stock available to them, and on the legal regime that determines whether and on what terms they may use it. End-to-end encryption, virtual private networks, and shielded payment networks each

function as instruments of privacy preservation whose deployment requires no legislative agreement, judicial intervention, or doctrinal evolution; they operate at the speed of adoption rather than the speed of governance, frequently orders of magnitude faster than the corresponding statutory or doctrinal correction. The presumptive legal availability of such technology is not itself the opacity balance, but a load-bearing condition for that balance to remain reachable by ordinary means. What this load-bearing role looks like in practice depends on which technologies are at issue, what they materially do for the citizens who adopt them, and how the state has positioned itself toward each. In this, the institutional specifics must be considered, both because the constitutional claim of the next section rests on them and because the technologies differ in ways that determine whether their availability genuinely preserves privacy against the countervailing incentives of the state or does so only nominally. We consider each technology in turn, for two reasons.

Privacy technology is distinctive in a way that makes it especially worth studying. Much of the privacy a citizen enjoys is governed by the state, through the constitutional protections, statutory limits, and procedural constraints that determine what state actors can see (in addition to also delimiting individuals' privacy as against private organizations, an important topic that is nonetheless separable from our analysis). Privacy technology is an important exception to this otherwise publicly-governed dynamic: in the modern era this technology is produced not by the state but by developers and firms,[4] and adopted directly by citizens, often at the speed of software and without explicit ex-ante institutional permission. This places a measure of privacy on a margin citizens can occupy for themselves, independent of whatever their institutions do or fail to do. This is what makes the state's posture toward it so consequential. Where privacy often directly depends on state protection, citizens who want more of it must persuade their institutions to provide it; where privacy can be adopted directly, the state's only move is whether to permit, restrict, or suppress what citizens are already doing. The opacity balance, in this domain, is set not by what the state provides but by what it allows citizens to provide for themselves, which is precisely why the terms on which it does so warrant close examination herein..

B. Privacy-Enhancing Technologies and Their Hostile Reception

The dynamic our theory describes is not only visible in the broad arc of surveillance law; it is visible in the contest between specific privacy-enhancing technologies and the government responses they provoke. Three such technologies have come to occupy distinctive positions in the contemporary opacity balance: end-to-end encryption of communications, virtual private networks, and shielded payment networks. They are not chosen for technological exhaustiveness

[4] Interestingly for our analysis, privacy technology like encryption was initially nearly exclusively the domain of the government and especially the military and national security apparatuses. Governments lost their monopoly on cryptographic dominance during the mid-1970s with the birth of public-key cryptography. The broader demographic and organizational shift, where private firms and citizens began producing and utilizing stronger, more widespread encryption than governments, solidified in the 1990s and early 2000s (Thompson and Park 2020).

but for what their reception reveals. Each constrains a different kind of legibility (the content of communications, the routing of network traffic, and the transactional history of payments, respectively) and each has diffused to citizens at the speed of software, outpacing the institutional channels that might have governed it. What follows traces what each technology does for the citizen who adopts it and how the U.S. government has responded to its spread. We take no position on whether any particular response reached the right result; our concern is with how the responses were produced, not whether they were correct. It is that procedural character, rather than the substantive merits, that bears on the constitutional argument of the sections that follow.

Consider first end-to-end encryption, the most widely adopted of the three and the one whose contest with the state is longest-running. End-to-end encryption (E2EE) describes a class of cryptographic protocols in which the content of a communication is visible only to the sender and the intended recipient. The intermediaries through which it passes, whether messaging platforms, network providers, or server operators, hold only ciphertext and lack the keys to decrypt it. Modern protocols such as the Signal Protocol, used in some form by Signal, WhatsApp, iMessage, and most major encrypted-messaging applications, generate a fresh key for each message rather than reusing one across a conversation, so that the compromise of any single key exposes only a small window of messages rather than the entire history, a property known as forward secrecy. The material consequence for the state is direct: an actor seeking the content of communications cannot obtain it through legal process directed at the platform, because the platform holds only encrypted data; access requires either reaching a participant's device or compelling a participant directly. Adoption has diffused rapidly since 2013, as encrypted messaging, default full-disk encryption on mobile devices, and Apple's introduction of Advanced Data Protection for iCloud in December 2022 extended E2EE to most of the everyday communications and stored data of ordinary users (Penney 2016).

The government's response has combined direct legal pressure, recurring legislative proposals, and sustained executive advocacy for "lawful access." The 2016 Apple v. FBI dispute over the San Bernardino iPhone, in which the government invoked the All Writs Act of 1789 to compel Apple to write and sign custom firmware defeating the device's own protections, ended without resolution when the FBI obtained access through a third-party vendor. Subsequent proposals have sought to condition platform liability protections on content scanning or to compel platforms to retain decryption capability; while none have been enacted, the cumulative posture has been one of steady pressure. That pressure runs against a settled technical consensus, articulated most influentially by a group of leading cryptographers, that "exceptional access" mechanisms such as backdoors, key escrow, and ghost participants, cannot be engineered to admit the state without degrading the security of encrypted communications for everyone (Abelson et al. 2015; Bellovin et al. 2013; Diffie and Landau 2007). The contest over encryption thus proceeds against the background of an inconvenient technical fact: the secure-with-access option the government has pursued does not exist.

The state's posture toward virtual private networks has been markedly more permissive, though that permissiveness is now under visible strain. A VPN establishes an encrypted tunnel between a user's device and a server operated by a provider, which forwards the user's traffic to its destination; to any observing internet service provider, network operator, or surveillance system positioned on the local network, only the tunnel is visible. To the services the user reaches, the traffic appears to originate from the provider rather than the user. The privacy this affords is real but structurally narrower than encryption's: the provider sees both ends of the user's traffic and becomes a single point of trust and of failure, which is why credible providers compete on "no-logs" policies, external audits, favorable jurisdiction, and diskless infrastructure. At the federal level the U.S. response has been notably accommodating with no licensing requirement, no criminal prohibition, no general treatment of VPN use as evidence of wrongdoing. Regulatory concern to date has been largely confined to the foreign ownership of particular providers and to subpoenas in specific investigations. But at the state level the picture is now contested: the wave of online age-verification regimes enacted since 2022 in Louisiana, Texas, Utah, Virginia, and elsewhere has driven a measurable rise in VPN use as residents route around platform geo-restrictions, prompting legislators in several states to weigh whether VPN use should itself be restricted, licensed, or treated as evidence of circumvention. No state has yet criminalized or licensed civilian use, but the trajectory of legislative interest alongside procurement bans on foreign-owned providers on government networks marks the federal permissiveness as an unstable equilibrium. Even so, the U.S. stance remains far more permissive than that of jurisdictions such as China, which licenses providers and bars unlicensed use, or Russia, which has blocked many providers outright. The federal permissiveness toward VPNs, set against the coercive response to shielded payments and the oscillating one to encryption, is itself part of the pattern: a stance that varies sharply across technologies, arrived at through no single deliberative process, and nowhere ratified as a considered public choice about how much network-layer legibility the polity is willing to forgo. This makes the state-level pressure on VPN access that is now surfacing worth watching as a test of whether the commitment holds.

The sharpest government response of the three privacy technologies we consider has been to shielded payment networks, where the state has moved most aggressively because the technology presses most directly against a legibility regime it has built up over half a century. Shielded payment networks are a heterogeneous set of cryptographic techniques that obscure the the linkages of sender, receiver, and amount that the public blockchains underlying most cryptocurrencies otherwise make permanently visible. Privacy-preserving cryptocurrencies build shielding into the protocol itself, through ring signatures and stealth addresses in Monero and zero-knowledge proofs in Zcash's shielded pool, whereas mixing services, or "tumblers," operate at a separate layer, pooling deposits from many users and permitting withdrawals to fresh addresses with no observable on-chain link. Among these mixers, the Tornado Cash protocol uses zero-knowledge proofs to let a depositor prove the right to withdraw without revealing which deposit was theirs. Materially, these tools defeat the blockchain-analysis tooling procured at scale by federal agencies on which financial-crime enforcement increasingly depends. These

mixing protocols' novelty is not financial privacy as such, which cash has always afforded, but the restoration of approximately cash-like privacy to a digital context whose default condition is comprehensive public disclosure.

The government's response has been correspondingly forceful. The Bank Secrecy Act of 1970 long ago made financial intermediaries surveil and report their customers' transactions, and the Financial Crimes Enforcement Network extended that framework to virtual-currency intermediaries, characterizing mixers as money transmitters in 2019 guidance (FinCEN 2019). In August 2022 the Office of Foreign Assets Control sanctioned the Tornado Cash protocol itself, marking the the first OFAC sanction of open-source software. OFAC added Tornado Cash's smart-contract addresses to the Specially Designated Nationals list, exposing any U.S. person's interaction with it to penalty. The sanction prompted litigation that culminated in the Fifth Circuit's holding in Van Loon v. Department of the Treasury (2024) that immutable smart contracts are not "property" reachable under the International Emergency Economic Powers Act, vacating the designation. Treasury proceeded to delist the protocol in March 2025. Separately, the federal prosecution of co-founder Roman Storm produced a partial conviction on a money-transmitting count in August 2025. The earlier 2018 conviction of Phantom Secure's chief executive for supplying encrypted devices had already set the pattern: the state's response to privacy infrastructure has centered on the developers who build it rather than on the users who misuse it.[5] What distinguishes the Tornado Cash episode is less the substantive question of financial privacy than the procedural character of the response that includes an executive sanction of open-source code, without statutory authorization specific to the practice, enforced against developers of a protocol they could not unilaterally alter subsequent to deployment on the network.

The substantive questions these technologies raise are genuinely difficult, requiring citizens and their representatives to weigh the protection of vulnerable groups and the interdiction of serious harms against the diffuse but cumulative erosion of privacy that broad surveillance imposes on the polity. We do not undertake to resolve that weighing, both because reasonable citizens will reach different conclusions and because their polity-specific resolution properly belongs to plenary democratic contestation rather than to academic analysis. What we do claim is that questions of this character that bear on both the substantive opacity balance and on the architecture through which all public sector commitments and actions are contested are "constitutional" questions that demand precisely the kind of transparent and comprehensive

[5] A contemporaneous European episode underscores that the dynamic is not a unique artifact of U.S. jurisprudence. Alexey Pertsev, another Tornado Cash co-founder, was convicted of money laundering by a Dutch court in May 2024 and sentenced to over five years; the conviction is under appeal. Proceeding under European anti-money-laundering law (through institutions wholly distinct from OFAC's authority, the All Writs Act, or the third-party doctrine), the prosecution nonetheless converged on the same target: the developer of a deployed, non-custodial protocol, held to answer for others' use of it. The offenses differ and we do not equate them; what generalizes is the structural response. That two legal orders as distinct as the United States and the Netherlands met the same protocol with the same move is evidence that the incentive we describe operates beneath the particularities of any single constitutional system.

contestation that the record just described too often forecloses. Across all three technologies the same sequence recurs: a privacy-preserving capability becomes available outside the channels of political contestation, citizens adopt it at the speed of software, and the state's response is more revelatory of its constitutional commitments than formal privacy doctrine alone. It is to that revealed pattern, and to what it reveals about the fidelity of a state's constitutional commitments in the digital era, that we now turn.

IV. The Opacity Balance as a Constitutional Dilemma

The pattern traced in the previous section, in which privacy-preserving capabilities diffuse to citizens and the state responds through ad hoc enforcement and administrative action, is a particular instance of the dynamic this paper has developed throughout: enforcement and administrative incentives cut consistently toward legibility, the corrections that would check them arrive late and asymmetrically, and past a sufficient rate of technological change, and especially as that rate accelerates, the result is a persistent drift no representative process authorized. That drift carries two implications, which this section develops. First, beyond a sufficient rate of change, it becomes a constitutional dilemma rather than a problem of ordinary regulation. Second, because it is a constitutional dilemma, the opacity balance should be settled through plenary democratic contestation, rather than left to the accretion of executive, regulatory, and enforcement decisions that fix a society's opacity in fact without ever putting the question to the polity.

The first implication follows from what makes the balance contested in the first place. The right level of opacity is not a contest between a privacy-loving citizenry and a legibility-hungry state, with the correct answer lying as far toward privacy as enforcement will allow. In democratic orders, the opacity balance is better understood as a contest among citizens themselves, who rationally value both sides of the ledger. The same person who wants the state not to read her messages can also want it to find an abducted child, recover her stolen car, interdict the financing of violence, and tax honestly the income others conceal, and each of these is a demand for legibility. Citizens do not hold a preference for privacy and a separate, weaker preference for security; they hold both at once, weigh them differently from one another, and weigh them differently across contexts. Some prefer a more empowered law enforcement, while others prefer a broader sphere of civil liberty, and most want some of each and disagree about the mix. Because legibility delivers public goods that citizens genuinely want, and privacy protects a sphere they genuinely value, neither can be maximized without sacrificing something the public actually demands. There is, in consequence, no technically correct opacity level that expertise could identify and install. Where the balance should sit is an irreducibly political question, in the precise sense that its answer depends on tradeoffs among competing legitimate preferences that only a representative process has the standing to make.

This is why the drift we have described is a constitutional matter. A regulatory lag is a gap between a settled public choice and its implementation, closed once institutions catch up. Some measure of regulatory lag is mechanical, in that technological and social change make public institutions lag this change by construction. What legibility drift produces is a movement of the opacity balance that no representative process ever chose, proceeding through administrative mechanisms that largely bypass public debate and legislative contestation. The balance still gets set; some configuration of preferences still prevails. But it is set by default, through the accumulated momentum of administrative and enforcement decisions each made for its own immediate purpose, rather than through any process that weighed the competing preferences of those who must live under it. The result is that a society's actual opacity comes to reflect not the balance its citizens would authorize but the balance that the actors operating its enforcement and administrative machinery find serviceable, making it a de facto answer to a question that was never put de jure.

When the same apparatus that benefits from legibility is also the one whose decisions fix the level of opacity, the arrangement has the form of a constitutional defect: the scope of the state's visibility into the governed is determined by the very actors that visibility empowers, with the governed's competing preferences never brought to bear. And because legibility obtained this way is not authorized by those it renders visible, the very capacity it is intended to enable can be self-defeating. Legibility is an input to state capacity, but so is legitimacy: a state secures cooperation, voluntary compliance, and trust in proportion as its oversight is seen to rest on consent rather than on accretion. Legibility extended faster than the polity will authorize can undermine the very capacity it was intended to obtain, trading a durable form of capacity grounded in legitimacy for a brittle one grounded in oversight enabled by technological change.

The remedy this implies is not a particular opacity level (which is itself contested within and across societies) but a particular process. If the balance is irreducibly political, then it ought to be set politically through legislation, adjudication, and public argument in which competing societal preferences are represented and the tradeoff is more fully made, rather than through the quiet accretion of administrative practice. To call the opacity balance constitutional is to say that it belongs to the polity as a whole and at the level of its basic commitments, not to whichever agency or vendor happens to field the next capability first. This is a claim about who decides and how, and it holds whatever substantive balance a given society would reach. A polity that, having deliberated, chooses extensive legibility for the security it buys has answered the question legitimately; a polity whose opacity drifts toward legibility through unexamined accretion has not answered it at all, however similar the two end states might look from outside.

Our second claim concerns urgency. That the opacity balance ought to be democratically legitimated is true at any rate of technological change; what acceleration does is make the failure to countenance it consequential and hard to reverse. When legibility advanced slowly, the gap between what the polity had authorized and what its apparatus had built could remain small, and contestation, arriving late, still had something to correct. As the rate of change rises, the gap

widens faster than contestation can close it, and because accumulated legibility is sticky, each unexamined increment narrows the range of balances that remain practically available to a polity that later tries to choose. Past a certain rate of change in information technology, acceleration thus converts a standing democratic obligation into an urgent one.[6] The longer the question goes unasked, the more the answer is settled by default, and the harder it becomes to ask in a form that any answer could still change. And as the institutional supply of privacy falls further behind, the private supply on which citizens increasingly rely becomes the more consequential battleground, which is why the state's posture toward privacy technology is not a peripheral question but the one on which the realized opacity balance increasingly turns.

The recent contest over automated license plate readers illustrates both the default-setting dynamic and the contestation that can interrupt it. A single vendor fielded a nationwide reader network across thousands of U.S. jurisdictions in under a decade, pooling billions of monthly plate reads into a searchable database, with no comprehensive statute authorizing the architecture and no moment at which any electorate was asked whether it wished to live under continuous, networked tracking of ordinary movement. For most of that period the balance was being set exactly as we describe: by the procurement decisions of local police and administrative bodies, each defensible according to the authorities that engaged in the adoption, yet nonetheless cumulatively fixing a level of locational transparency that no one subject to it had chosen. As the scale of the system became publicly visible, it provoked the contestation that had been bypassed, through investigative reporting, organized public objection, municipal decisions to decline or remove the cameras, and the first comprehensive state statutes regulating their use. The response did not establish that the cameras are impermissible; reasonable citizens differ on that, and that is our point. What it established is that the question is one a polity can and will engage once it is surfaced as a question, and that the failure was never the technology's existence but the years in which its adoption outran the public's opportunity to weigh it. The episode demonstrates the deliberative apparatus is not incapable of governing legibility, only chronically behind it, and that the remedy we describe is available wherever the constitutional question is posed in time to matter.

V. Conclusion

We have argued that privacy and legibility are both necessary conditions of a capable democratic state, and a society's observed level of effective non-visibility is better called opacity. Any observed opacity level is the unstable product of a standing institutional dilemma: the same apparatus charged with protecting privacy is the one whose enforcement and administrative actors benefit most from legibility, so the supply of opacity is structurally precarious even before

[6] To say that we are entering a phase of accelerating change in information technology considering the advent of artificial intelligence may be to understate the magnitude, which makes contestation over the use of artificial intelligence for government administration and enforcement a contemporary margin by which to consider our arguments herein.

technology enters. Technology then enters as a sequence of shocks, and the dilemma turns dynamic. Legibility-enhancing capabilities arrive fast and cheap; the corrections that would restore a representative balance arrive slowly, reach only forward, and leave the accumulated stock of visibility in place. The result is a legibility ratchet, and past a sufficient rate of change an accelerating one, in which opacity drifts toward legibility through channels no representative process authorized. We then illustrate this dynamic through the contest between specific privacy technologies and the U.S. government responses they provoke. More deeply, our modern privacy technology examples suggest that where the opacity balance should sit is an irreducibly political question, one that ought to be settled through plenary contestation but that, at present, is often settled by the default of administration and enforcement in practice.

It remains to say what a state's conduct in this domain reveals, and why it is worth watching beyond its immediate stakes. The argument rests on an uncontroversial comparative premise: states differ, and can be ranked, in how faithfully their institutions honor the commitments their constitutional orders profess. This is the shared foundation of a large body of practice and scholarship, from rule-of-law indices that score the gap between professed and realized constraint, to human-rights monitoring that assesses how fully ratified guarantees are delivered, to the literature distinguishing genuine constitutional constraint from the merely nominal or "sham" constitution whose written guarantees bind nothing (Law and Versteeg 2024). The common premise is simply that a constitution is a set of commitments meant to bind the exercise of concentrated authority, that the rules of a given society succeed to varying and observable degrees in actually binding it, and that some societies genuinely shackle the power they organize while others profess the same constraints and live under far less of them.

The commitments whose fidelity is most informative are those that some identifiable class of state actors is most reliably driven to breach. A commitment no state actor has reason to violate is honored under almost any arrangement, and so cannot distinguish binding constraints from decorative ones; a commitment that powerful actors face consistent incentives to override is honored only where institutions actually check that pressure. Privacy is such a commitment. The actors who wield the state's investigative, prosecutorial, and administrative authority face standing incentives to expand what the state can see, arising from citizen demands for protection, from the pursuit of their own authorized objectives, and from the simple utility of legibility to governing. These incentives are not a defect of particular officials; they are a structural feature of the enforcement function, present in some form in every state. Privacy technology is the domain in which they press most consistently, because it places citizen activity directly beyond the reach those actors predictably seek to extend. A state's stance toward it therefore lands where the constitutional commitment to privacy is hardest to keep, and where institutional fidelity, if present, is most costly and so most revealing. Constitutionalism’s extent, on this view, can be examined through the set of arrangements that hold constraints in place where the pressure against them is greatest; a constitutional order that constrained only conduct no government actor was incentivized to engage in would constrain nothing of consequence.

The citizen's side makes this area of government response even more revelatory when it comes to privacy technology. Because this technology is a good citizens can supply to themselves, the state's response is not a choice whether to provide privacy but whether to permit one that firms produce and citizens are already exercising. A state that merely fails to supply privacy might lack the capacity or the occasion; a state that moves to close a channel of self-protection citizens have adopted on their own must affirmatively reach into a margin of choice already open to them, and what it does there reveals whether its professed commitment survives contact with a citizenry actually exercising it.

Because the domain is so demanding, what a state's conduct there reveals carries beyond privacy. If an arrangement checks the legibility push where that push is strongest, that is suggestive evidence its constraints bind more generally; if its authorities give way precisely where the push is hardest to resist, the question follows outward, since an order whose commitments yield under the greatest pressure invites doubt about their hold where the pressure is merely ordinary. A state's conduct toward privacy technology is in this sense more than a privacy story. It is among the more searching tests of whether a constitutional order's commitments bind in fact or only on paper, an arguable bellwether for the fidelity of the whole. And for a state whose privacy commitments are already weak, the stakes are not merely diagnostic: where institutional protection has failed, citizen-adopted technology is the last remaining source of privacy, so a move to suppress it does not just reveal the weakness of the commitment but extinguishes the substitute that weakness had made more important for preserving what privacy individuals could.

That fidelity has two margins. The first is substantive: whether a state, confronted with a capability its enforcement apparatus has every reason to suppress, leaves the corresponding sphere of privacy presumptively intact or moves to close it. The second is procedural: how it does so. A state that constrains privacy technology through transparent, legislatively authorized processes honors the deliberative architecture through which constitutional commitments are meant to be set, even where the outcome is restrictive; one that proceeds through opaque executive sanction, prosecution of developers without antecedent legislative authority, or unilateral pressure on intermediaries reveals something about its fidelity to that architecture independent of any result. The U.S. record exhibits both registers and resists a clean characterization: broadly permissive toward consumer encryption, sharply coercive toward shielded payment networks, ambivalent and increasingly contested toward the VPN tools between them.

We further stress that a permissive stance does not by itself certify that a state has found the right balance, any more than a restrictive one condemns it; reasonable democracies situate the balance differently, as the divergence among the United States' treatment of consumer encryption, the United Kingdom's notice powers under the Online Safety Act, and Australia's industry-assistance regime attests. Nor is such a stance informative where permissiveness reflects incapacity rather than constraint: an arrangement that tolerates a technology no actor within it could suppress

reveals nothing about whether its constraints bind. For the formal posture of a government toward privacy to be revelatory, this would therefore require a state whose actors demonstrably could bring its coercive power to bear and whose institutions must decide whether they will. Only against such a backdrop of public institutional definition through transparent procedural and substantive constitutional guarantees, are the domains in which the state stays its hand legible as the legitimate output of institutions that constrained it.

Importantly, though, the structural dilemma is not in every dimension zero-sum. The technologies we have traced cast the tension in its starkest form, each subtracting from the state's view so that every gain to privacy is a loss to legibility. But an emerging class of technology, privacy-preserving computation, relaxes that tension along certain margins. Selective-disclosure tools built on zero-knowledge proofs let a party prove a specific claim without revealing anything further (Alston and Cossar 2026); differential privacy permits accurate inference about a population while bounding what can be learned about any individual; fully homomorphic encryption permits computation over data that is never decrypted. Each lets a verifier learn the particular fact its authority extends to while leaving everything else unseen, and so each mitigates the very tradeoff the static dilemma treats as fixed. They are evidence that the frontier between privacy and legibility is not a fixed line along which one side's gain is always the other's loss, but one that technological ingenuity can push outward, so that legitimate demands for verification and for privacy need not always be satisfied at one another's expense.

The promise is real but bounded, in two ways worth naming so it is not mistaken for a solution. The first is functional: these are tools of verification, of proving a discrete fact, and much of the activity that increases the legibility drift is not verification at all, so for those domains the tradeoff remains as stark as before. The second is deeper, and returns us to the argument's center. These technologies answer how a fact may be confirmed without wholesale disclosure; they are silent on the prior and properly democratic question of what the state should be entitled to confirm at all. By making each act of verification feel less intrusive, these selective disclosure mechanisms could lower the political resistance that would otherwise force that question to be asked, licensing a quiet expansion of legibility into domains a polity might, on reflection, have chosen to keep opaque. A tool that lowers the cost of verification can weaken the check that ought to govern its scope. These technologies are most valuable, then, exactly where they inform democratic contestation rather than substitute for it, which is the note on which our whole argument insists. The balance between privacy and legibility is a constitutional question; technology can enlarge the menu of answers, and can reveal how faithfully a state's institutions honor the answer it professes, but it cannot relieve a democratic polity of the obligation to choose.

Sources Cited:

Court cases and statutes

Carpenter v. United States, 585 U.S. 296 (2018).

Chatrie v. United States, No. 25-112 (U.S., cert. granted Jan. 16, 2026; argued Apr. 27, 2026). On review of United States v. Chatrie, 107 F.4th 319 (4th Cir. 2024), reh'g en banc granted, opinion vacated, judgment aff'd, 2025 WL 1242063 (4th Cir. 2025) (en banc). District court below: United States v. Chatrie, 590 F. Supp. 3d 901 (E.D. Va. 2022).

Clapper v. Amnesty International USA, 568 U.S. 398 (2013).

In the Matter of the Search of an Apple iPhone (In re Apple), No. ED 15-0451M (C.D. Cal. Feb. 16, 2016) (order compelling Apple to assist).

Katz v. United States, 389 U.S. 347 (1967).

NAACP v. Alabama ex rel. Patterson, 357 U.S. 449 (1958).

Riley v. California, 573 U.S. 373 (2014).

Smith v. Maryland, 442 U.S. 735 (1979).

State v. Loomis, 2016 WI 68, 371 Wis. 2d 235, 881 N.W.2d 749 (2016), cert. denied, 137 S. Ct. 2290 (2017).

United States v. Jones, 565 U.S. 400 (2012).

United States v. Maynard, 615 F.3d 544 (D.C. Cir. 2010).

United States v. Miller, 425 U.S. 435 (1976).

Van Loon v. Department of the Treasury, 122 F.4th 549 (5th Cir. 2024).

All Writs Act, 28 U.S.C. § 1651 (originally Judiciary Act of 1789, § 14, 1 Stat. 81–82).

Assistance and Access Act (2018) Telecommunications and Other Legislation Amendment (Assistance and Access) Act 2018 (Cth), Act No. 148 of 2018 (Australia).

Bank Secrecy Act of 1970, Pub. L. No. 91–508, 84 Stat. 1114 (1970).

EARN IT Act (2020) Eliminating Abusive and Rampant Neglect of Interactive Technologies Act, S. 3398, 116th Cong.; reintroduced as S. 3538, 117th Cong. (2022); S. 1207, 118th Cong. (2023).

Lawful Access to Encrypted Data Act (2020) S. 4051, 116th Cong.

Online Safety Act (2023) Online Safety Act 2023, c. 50 (UK), §§ 121–122.

Oregon SB 1516 (2026) Senate Bill 1516, 83rd Oregon Legislative Assembly, 2026.

USA FREEDOM Act (2015) Pub. L. No. 114-23, 129 Stat. 268.

USA PATRIOT Act (2001) Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism Act of 2001, Pub. L. No. 107-56, 115 Stat. 272.

Scholarly works

Abelson, H., Anderson, R., Bellovin, S. M., Benaloh, J., Blaze, M., Diffie, W., Gilmore, J., Green, M., Landau, S., Neumann, P. G., Rivest, R. L., Schiller, J. I., Schneier, B., Specter, M. A. and Weitzner, D. J. (2015) 'Keys under doormats: Mandating insecurity by requiring government access to all data and communications', Journal of Cybersecurity, 1(1), pp. 69–79.

Acemoglu, D. and Robinson, J. A. (2012) Why Nations Fail: The Origins of Power, Prosperity, and Poverty. New York: Crown Publishers.

Acquisti, A., Brandimarte, L. and Loewenstein, G. (2015) 'Privacy and human behavior in the age of information', Science, 347(6221), pp. 509–514.

Acquisti, A., Taylor, C. and Wagman, L., 2016. The economics of privacy. Journal of economic Literature, 54(2), pp.442-492.

Allen, A. L. (2011) Unpopular Privacy: What Must We Hide?. New York: Oxford University Press.

Alston, E. and Cossar, S. (2026) 'The institutional privacy dilemma: Technological change and the coevolution of privacy institutions', working paper.

Babazadeh, N. (2018) 'Concealing evidence: "Parallel construction," federal investigations, and the Constitution', Virginia Journal of Law & Technology, 22(1), pp. 1–68.

Bamberger, K. A., Canetti, R., Goldwasser, S., Wexler, R. and Zimmerman, E. J. (2022) 'Verification dilemmas in law and the promise of zero-knowledge proofs', Berkeley Technology Law Journal, 37(1), pp. 1–70.

Becker, S. O., Pfaff, S., Hsiao, Y. and Rubin, J. (2023) 'Competing social influence in contested diffusion: Luther, Erasmus and the spread of the Protestant Reformation', ESI Working Paper 23-03.

Bellovin, S. M., Blaze, M., Clark, S. and Landau, S. (2013) 'Going bright: Wiretapping without weakening communications infrastructure', IEEE Security & Privacy, 11(1), pp. 62–72.

Benati, G., Carugati, F. and Léon-Ablán, G., 2026. Writing and the State: Information and State Capacity in Mesopotamia. Working paper.

Besley, T. and Persson, T. (2009) 'The origins of state capacity: Property rights, taxation, and politics', American Economic Review, 99(4), pp. 1218–1244.

Black, J. (1987) The English Press in the Eighteenth Century. London: Croom Helm.

Bloch-Wehba, H. (2021) 'Transparency's AI problem', in Data and Democracy, Knight First Amendment Institute at Columbia University, Essay Series.

Boerner, L., Rubin, J. and Severgnini, B. (2021) 'A time to print, a time to reform', European Economic Review, 138, 103826.

Bowles, J. (2024) 'Identifying the rich: Registration, taxation, and access to the state in Tanzania', American Political Science Review, 118(2), pp. 602–618.

Brambor, T., Goenaga, A., Lindvall, J. and Teorell, J. (2020) 'The lay of the land: Information capacity and the modern state', Comparative Political Studies, 53(2), pp. 175–213.

Cheng, C., Stasavage, D. and Wang, Y. (2023) 'The written word and development of the state in China and Europe', Research Group on Political Institutions and Economic Policy (PIEP).

Childe, V. G. (1936) Man Makes Himself. London: Watts & Company.

Childe, V. G. (1950) 'The urban revolution', The Town Planning Review, 21(1), pp. 3–17.

Christensen, D. and Garfias, F. (2021) 'The politics of property taxation: Fiscal infrastructure and electoral incentives in Brazil', Journal of Politics, 83(4), pp. 1399–1416.

Cingolani, L. (2023) 'Infrastructural state capacity in the digital age: What drives the performance of COVID-19 tracing apps?', Governance, 36(3), pp. 833–853.

Cohen, J. E. (2013) 'What privacy is for', Harvard Law Review, 126(7), pp. 1904–1933.

Collingridge, D. (1980) The Social Control of Technology. New York: St. Martin's Press.

D'Arcy, M. and Nistotskaya, M. (2017) 'State first, then democracy: Using cadastral records to explain governmental performance in public goods provision', Governance, 30(2), pp. 193–209.

Diffie, W. and Landau, S. (2007) Privacy on the Line: The Politics of Wiretapping and Encryption, updated and expanded edn. Cambridge, MA: MIT Press.

Dittmar, J. E. (2011) 'Information technology and economic change: The impact of the printing press', The Quarterly Journal of Economics, 126(3), pp. 1133–1172.

Donohue, L. K. (2016) The Future of Foreign Intelligence: Privacy and Surveillance in a Digital Age. New York: Oxford University Press.

Downes, L. (2009) The Laws of Disruption: Harnessing the New Forces that Govern Life and Business in the Digital Age. New York: Basic Books.

Etzioni, A. (1999) The Limits of Privacy. New York: Basic Books.

Foucault, M. (1975) Discipline and Punish: The Birth of the Prison, trans. A. Sheridan. New York: Pantheon Books (1977).

Gadd, I. and Gillespie, A. (eds.) (2006) Print and Power in Early Modern Europe (1500–1800). Cambridge: Cambridge University Press.

Gavison, R. (1980) 'Privacy and the limits of law', Yale Law Journal, 89(3), pp. 421–471.

Goitein, E. and Patel, F. (2015) What Went Wrong with the FISA Court. New York: Brennan Center for Justice.

Goody, J. (1986) The Logic of Writing and the Organization of Society. Cambridge: Cambridge University Press.

Habermas, J. (1996) Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy, trans. W. Rehg. Cambridge, MA: MIT Press.

Hermalin, B. E. and Katz, M. L. (2006) 'Privacy, property rights and efficiency: The economics of privacy as secrecy', Quantitative Marketing and Economics, 4(3), pp. 209–239.

Hirshleifer, J. (1980) 'Privacy: Its origin, function, and future', Journal of Legal Studies, 9(4), pp. 649–664.

Hudson, M. (2000) 'Mesopotamia and classical antiquity', American Journal of Economics and Sociology, 59, pp. 3–25.

Law, D. S. and Versteeg, M. (2013) 'Sham constitutions', California Law Review, 101(4), pp. 863–952.

Lee, M. M. and Zhang, N. (2017) 'Legibility and the informational foundations of state capacity', Journal of Politics, 79(1), pp. 118–132.

Lessig, L. (2006) Code: Version 2.0. New York: Basic Books.

Levi, M. (1988) Of Rule and Revenue. Berkeley, CA: University of California Press.

Liverani, M. (2006) Uruk: The First City. London: Equinox.

Logan, W. A. (2009) Knowledge as Power: Criminal Registration and Community Notification Laws in America. Stanford, CA: Stanford University Press.

Logan, W. A. (2011) 'Megan's laws as a case study in political stasis', Wake Forest Law Review, 46, pp. 167–207.

Mann, M. (1984) 'The autonomous power of the state: Its origins, mechanisms and results', European Journal of Sociology, 25(2), pp. 185–213.

Mansbridge, J. (2003) 'Rethinking representation', American Political Science Review, 97(4), pp. 515–528.

Marthews, A. and Tucker, C. (2017) 'Government surveillance and internet search behavior', MIT Sloan Working Paper / SSRN.

Miller, C. R. (1983) Technical and Cultural Prerequisites for the Invention of Printing in China and the West. San Francisco: Chinese Materials Center.

Mukerji, C. (2011) 'Jurisdiction, inscription, and state formation: Administrative technologies in early modern France'.

Murphy, R. S. (1996) 'Property rights in personal information: An economic defense of privacy', Georgetown Law Journal, 84(7), pp. 2381–2417.

Nissen, H. J. (1993) 'The context of the emergence of writing in Mesopotamia and Iran', in J. Curtis (ed.) Early Mesopotamia and Iran: Contact and Conflict 3500–1600 BC. London: The British Museum Press, pp. 54–71.

Nissenbaum, H. (2010) Privacy in Context: Technology, Policy, and the Integrity of Social Life. Stanford, CA: Stanford University Press.

North, D. C. and Weingast, B. R. (1989) 'Constitutions and commitment: The evolution of institutions governing public choice in seventeenth-century England', Journal of Economic History, 49(4), pp. 803–832.

North, D. C., Wallis, J. J. and Weingast, B. R. (2009) Violence and Social Orders: A Conceptual Framework for Interpreting Recorded Human History. Cambridge: Cambridge University Press.

Penney, J. W. (2016) 'Chilling effects: Online surveillance and Wikipedia use', Berkeley Technology Law Journal, 31(1), pp. 117–182.

Pettit, P. (1997) Republicanism: A Theory of Freedom and Government. Oxford: Oxford University Press.

Posner, R. A. (1978) 'The right of privacy', Georgia Law Review, 12(3), pp. 393–422.

Raymond, J. (1996) The Press and Popular Political Opinion in England 1660–1714.

Reidenberg, J. R. (1998) 'Lex informatica: The formulation of information policy rules through technology', Texas Law Review, 76(3), pp. 553–593.

Richards, N. M. (2013) 'The dangers of surveillance', Harvard Law Review, 126(7), pp. 1934–1965.

Richards, N. M. (2015) Intellectual Privacy: Rethinking Civil Liberties in the Digital Age. New York: Oxford University Press.

Richardson, S. (2017) 'Messaging and the gods in Mesopotamia: Signals and systematics', in R. Talbert and F. Naiden (eds.) Mercury's Wings: Exploring Modes of Communication in the Ancient World. New York: Oxford University Press, pp. 167–194.

Roberts, A. (2015) 'A republican account of the value of privacy', European Journal of Political Theory, 14(3), pp. 320–344.

Rogowski, J.C., Gerring, J., Maguire, M. and Cojocaru, L., 2022. Public infrastructure and economic development: Evidence from postal systems. American Journal of Political Science, 66(4), pp.885-901.

Rubin, J. (2014) 'Printing and Protestants: An empirical test of the role of printing in the Reformation', Review of Economics and Statistics, 96(2), pp. 270–286.

Rubin, J. (2017) Rulers, Religion, and Riches: Why the West Got Rich and the Middle East Did Not. New York: Cambridge University Press.

Scott, J. C. (1985) Weapons of the Weak: Everyday Forms of Peasant Resistance. New Haven, CT: Yale University Press.

Scott, J. C. (1998) Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed. New Haven, CT: Yale University Press.

Scott, J. C. (2009) The Art of Not Being Governed: An Anarchist History of Upland Southeast Asia. New Haven, CT: Yale University Press.

Selinger, E. and Hartzog, W. (2019) 'The inconsentability of facial surveillance', Loyola Law Review, 66(1), pp. 101–122.

Slobogin, C. (2008) 'Government data mining and the Fourth Amendment', University of Chicago Law Review, 75(1), pp. 317–341.

Soifer, H. (2008) 'State infrastructural power: Approaches to conceptualization and measurement', Studies in Comparative International Development, 43(3–4), pp. 231–251.

Solove, D. J. (2006) 'A taxonomy of privacy', University of Pennsylvania Law Review, 154(3), pp. 477–564.

Solove, D. J. (2008) Understanding Privacy. Cambridge, MA: Harvard University Press.

Solove, D. J. (2021) 'The myth of the privacy paradox', George Washington Law Review, 89(1), pp. 1–51.

Stasavage, D. (2021) 'Biogeography, writing, and the origins of the state', in A. Bisin and G. Federico (eds.) The Handbook of Historical Economics. London: Academic Press.

Stigler, G. J. (1980) 'An introduction to privacy in economics and politics', Journal of Legal Studies, 9(4), pp. 623–644.

Tilly, C. (1990) Coercion, Capital, and European States, AD 990–1990. Cambridge, MA: Blackwell.

Vladeck, S. I. (2014) 'Standing and secret surveillance', I/S: A Journal of Law and Policy for the Information Society, 10(2), pp. 551–576.

Vladeck, S. I. (2015) 'The FISA Court and Article III', Washington & Lee Law Review, 72(3), pp. 1161–1208.

Warren, S. D. and Brandeis, L. D. (1890) 'The right to privacy', Harvard Law Review, 4(5), pp. 193–220.

Westin, A. F. (1967) Privacy and Freedom. New York: Atheneum.

Yoffee, N. (2005) Myths of the Archaic State: Evolution of the Earliest Cities, States, and Civilizations. Cambridge: Cambridge University Press.

Zuboff, S. (2019) The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power. New York: PublicAffairs.

Government documents and reports

Administration White Paper (2013) Bulk Collection of Telephony Metadata Under Section 215 of the USA PATRIOT Act. Washington, DC: U.S. Government, 9 August.

DOJ (2015) Department of Justice Policy Guidance: Use of Cell-Site Simulator Technology, U.S. Department of Justice, 3 September.

DOJ-SDCA (2018) 'CEO of Encrypted Communications Company Pleads Guilty', U.S. Attorney's Office, Southern District of California, Press Release, 2 October.

ODNI (2019) Statistical Transparency Report Regarding the Use of National Security Authorities — Calendar Year 2018. Office of the Director of National Intelligence, April.

OFAC (2022) 'Treasury sanctions notorious virtual currency mixer Tornado Cash', U.S. Department of the Treasury, Office of Foreign Assets Control, Press Release, 8 August.

PCLOB (2020) Report on the Government's Use of the Call Detail Records Program Under the USA FREEDOM Act. Privacy and Civil Liberties Oversight Board, Washington, DC.

USAO-SDNY (2025) 'Founder of Tornado Cash crypto mixing service convicted of knowingly transmitting criminal proceeds', U.S. Attorney's Office, Southern District of New York, Press Release, 6 August.

U.S. Treasury (2025) 'Tornado Cash delisting', U.S. Department of the Treasury, Press Release, 21 March.

Policy, Journalism and Industry Sources

ACLU of Massachusetts (2025) Get The FLOCK Out: Resource Guide.

Apple Inc. (2022) 'Apple advances user security with powerful new data protections', Apple Newsroom, 7 December.

Garvie, C., Bedoya, A. M. and Frankle, J. (2016) The Perpetual Line-Up: Unregulated Police Face Recognition in America. Washington, DC: Center on Privacy & Technology at Georgetown Law.

Google (2021) Supplemental Information on Geofence Warrants in the United States. Google LLC.

Greenwald, G. (2013) 'NSA collecting phone records of millions of Verizon customers daily', The Guardian, 6 June.

Hill, K. (2020) 'The secretive company that might end privacy as we know it', The New York Times, 18 January.

Nangueneri, S. (2026) 'A new Oregon law regulates police use of license plate readers', Oregon Capital Chronicle, 23 April.

NBC News (2025) 'Flock police cameras scan billions per month, sparking growing protests', NBC News, 1 November.

Oregon Public Broadcasting (2026) 'Think Out Loud: Oregon State Police sharing data with ICE', OPB, 11 May.

Wang, N., McDonald, A., Bateyko, D. and Tucker, E. (2022) American Dragnet: Data-Driven Deportation in the 21st Century. Washington, DC: Center on Privacy & Technology at Georgetown Law.

Wilson Thompson, A. and Park, C. (2020) Privacy's Best Friend: The Importance of Encryption in Protecting Consumer Privacy. Washington, DC: New America's Open Technology Institute, 24 August.