

Privacy-Preserving Age Verification—and Its Limitations

Steven M. Bellovin *

smb476@georgetown.edu

<https://www.cs.columbia.edu/~smb>

October 2025

Abstract

To a first technical approximation, it is straightforward to construct a privacy-preserving, credential-based age verification system for the Worldwide Web. However, the legal, economic, and social obstacles are formidable, and possibly insurmountable, especially in certain countries.

1 Introduction

Many jurisdictions around the world have imposed age verification requirements for access to certain content, e.g., pornography. These include the United Kingdom [35, 36], the European Union [14], and a number of U.S. states, including among others Texas [34] and Louisiana [25]. Australia has passed a law barring children under 16 from social media sites [6]. In the U.S., the constitutionality of such laws for accessing adult content was recently upheld by the Supreme Court [15]; the status of age restrictions for other purposes is still being litigated [5].

The problem is how to implement such schemes, especially if you want continuous verification.

There are several approaches described in [4], including self-attestation (the user states how old they are), biometric-based evaluation, and credential-based authentication. The first is obviously security theater and not worth pursuing. The second may or may not be possible in principle (though I'm quite skeptical), but as the authors note, the known algorithms do poorly near legally significant boundaries, e.g., under or over 18. Furthermore, they point out that [4, p. 29] “While some companies that provide age estimation based on face images are willing to participate NIST’s FATE-AEV program, effectively none provide open-source code or training data that would allow audits by academic researchers or humans rights advocates. This makes it impossible to tell to what extent self-reported or NIST correctness metrics actually correspond to performance on the wide variety of demographics, image variations, and other real world complications that challenge deployed age estimators.” Given the known causes of error in facial recognition, including race and gender [27], it is reasonable to assume that age recognition algorithms would suffer analogous flaws and more.

Google is rolling out a behavioral-based system, probably using machine learning [9]. All they’ve disclosed about how it works is “Our age estimation model uses machine learning to interpret a variety of signals already associated with a user’s account, such as the types of information a user has searched for or the categories of videos they’ve watched on YouTube. These signals help us determine whether a user is likely over or under the age of 18.” It seems likely that this scheme would have similar boundary issues. Consider one obvious signal: searches for college admission information. Certainly, most people performing such searches in the U.S. would be high school seniors, aged 17 or 18. But non-traditional students or those who have to do mandatory military service before college (e.g., most Israelis) might be considerably older when applying to college, and those who have skipped a grade could easily be younger—and 18 years old is a common legal boundary. They propose handling incorrect under-18 values via an appeal process: “If we incorrectly estimate a user to be under 18, the user has the option to correct their age, including by uploading a photo of their government ID or a selfie.” But selfies are just a form of biometric, with all that implies for accuracy and privacy. (The other problem, of course, is someone creating two Google accounts, and using one only to appear older. Without more knowledge of just what Google will look at, that might of course be harder than it sounds.)

*Senior Affiliate Scholar, Georgetown University Institute of Technology Law & Policy; Percy K. and Vida L.W. Hudson Professor Emeritus of Computer Science, Columbia University.

We're thus left with credentials. In some sense, it's easy: have the user present a credential showing their age. But that raises several issues, including the authenticity of the credential, verification that the user of the credential is in fact the legitimate possessor of it, and of course privacy. Note that privacy includes not just use or resale of whatever data a user provides to verify their age, but also the security of the verifying company's database, especially if there is a retention requirement to permit audits by government agencies. Furthermore, if age verification in some jurisdiction is necessary only for access to sexual content, the mere existence of a person's record in the verifier's database is equivalent to statement that that person wishes to view pornography. This could be harmful to many relationships.

There are strong technical solutions to some of these issues, especially authenticity and privacy; however, these raise other issues that in some jurisdictions may prove insurmountable.

The basic concept is to use a privacy-preserving credential scheme, and in particular one that can carry other attributes such as age assertions. One such scheme was proposed by Camenisch and Lysyanskaya [10]; an implementation in a different context, deterring intimate image abuse, was done by Gorman et al. [16] and a comprehensive U.S. legal analysis of the use of the scheme (including a more accessible description of the protocol) was done by Zhang and Bellovin [41].

There are, however, several difficult problems, including availability of Camenisch-Lysyanskaya (CL) credentials for all users, and economics: who should pay the cost of the necessary infrastructure? From a technical perspective, a key security assumption behind the CL protocol only holds if these credentials are used for multiple purposes. Furthermore, the international nature of the Web presents its own set of governance difficulties.

In the remainder of this note, I describe the CL protocol, how to use it for the web, and the challenges to its use. I note that many of the obstacles will apply to *any* age verification scheme.

2 The Ideal Technical Solution

2.1 The Camenisch-Lysyanskaya Protocol

At its highest level, the CL protocol is simple. A site known as an *Identity Provider (IDP)* issues what is called a *primary credential* by Zhang and Bellovin [41]. During this process, the IDP can ask for any sort of information it wishes; for this purpose, proof of age is most important.

The possessor of a primary credential (and its associated private key, of course) can ask the IDP for any number of *subcredentials*; these subcredentials can be used to log in to any site that speaks this protocol using zero-knowledge proofs. The subcredentials have three crucial properties: they are provably derivable from a primary credential issued by a trusted IDP; they cannot be linked to each other; and they cannot be linked to a primary credential. Thus, whoever accepts them is assured that they're valid, but does not know who the possessor is.

There are two optional extensions described by Camenisch and Lysyanskaya. One, which is not relevant here, provides for *revocable anonymity*: a *deanonymization agent* can decrypt an encrypted version of an identifier known to the IDP, which has presumably kept a record of which user is associated with which identifier. The other extension is more interesting to us: the subcredentials can carry a series of binary attributes such as "over 18," "over 21," etc. The integrity of these, too, are covered by the proof of validity of the subcredential.

2.2 Using the CL Protocol on the Web

In some sense, using the CL protocol on the Web is straightforward: simply implement the necessary code in all browsers and Web servers. The problem with such a simplistic answer is obvious: there are probably billions of browsers and at least hundreds of millions of Web servers; updating them all is somewhere between extremely difficult and impossible, and that's without even considering the time to update the TLS specification and its common implementations, e.g., OpenSSL.

Instead, I suggest using an existing TLS mechanism: client-side certificates [30, §4.3.2]. Without going into detail, mechanisms akin to those described in an obsolete Internet draft [7] can be used to obtain certificates: use the CL protocol with a subcredential, via a browser extension, to login to a certificate authority (CA), which will return a standard X.509 certificate that can be used with any Web browser and server. Extension fields in the certificate can be used to carry age assertions. For users who have multiple devices, there may be built-in platform-specific key transport mechanism such as Apple's Keychain; alternatively, schemes such as described by Koh, Bellovin, and Nieh [21] could be used. It may be easier (and more privacy-preserving) to obtain a separate certificate for each device (but this has its own disadvantages, as discussed below).

Much of this can be automated. Once a user has applied for a primary credential, it and the associated private key can be stored locally by their browser. Some number of subcredentials can be obtained and cached at the same time. A site that wishes to verify a user's age would send a CertificateRequest message in its TLS negotiation; at that point, the browser would contact a CA to obtain a certificate and then continue the TLS login process. Naturally, advance-use certificates can be cached as well.

There are some disadvantages to using a separate CA. First, it is another entity that must be funded, adding to the cost of the scheme. Second, and quite crucially, the CAs *must* be organizationally separate from the Identity Provider; otherwise, there is too much information known to one party, thereby potentially compromising user privacy. Furthermore, communications with the CA and probably the IDP should be done over Tor [13], to prevent linkage via IP address. Finally, if there are no cached advance-use certificates, this would add a noticeable delay during first access to a site.

3 Insurmountable Obstacles

There are a number of insurmountable obstacles to using CL or any similar solution.

Obtaining a Primary Credential: The biggest single problem with using CL is access to and authentication by IDPs. This is especially serious in countries like the U.S. and the U.K., where there is no national ID card. Other factors that hurt access are age, poverty, distance from an Identity Provider or the government agency that would issue the necessary ID, and more. A more complete analysis is given in §V.C of [41]; most of that section would apply to many countries and not just the United States. Two quotes, from a U.S. Supreme Court opinion involving a state government requirement for a government-issued ID for voting, spells out many of the issues [12]:

Both evidence in the record and facts of which we may take judicial notice, however, indicate that a somewhat heavier burden may be placed on a limited number of persons. They include elderly persons born out of State, who may have difficulty obtaining a birth certificate; persons who because of economic or other personal limitations may find it difficult either to secure a copy of their birth certificate or to assemble the other required documentation to obtain a state-issued identification; homeless persons; and persons with a religious objection to being photographed.

and

The first set of burdens shown in these cases is the travel costs and fees necessary to get one of the limited variety of federal or state photo identifications needed to cast a regular ballot under the Voter ID Law. The travel is required for the personal visit to a license branch of the Indiana Bureau of Motor Vehicles (BMV), which is demanded of anyone applying for a driver's license or nondriver photo identification. The need to travel to a BMV branch will affect voters according to their circumstances, with the average person probably viewing it as nothing more than an inconvenience. Poor, old, and disabled voters who do not drive a car, however, may find the trip prohibitive, witness the fact that the BMV has far fewer license branches in each county than there are voting precincts.

Many of these issues would apply in other countries with large, sparsely populated areas besides the U.S.; Australia, Brazil, Canada, China, and Russia come to mind, as well as much of North Africa.

The fact that multiple forms of ID are acceptable, while solving some problems, in fact exacerbates the fraud issue. A person may have a driver's license from more than one state or province (which shouldn't happen but to my certain knowledge does), multiple passports (when I worked for the U.S. government, I had both a personal and an official passport), odder forms of ID (Texas permits use of a state handgun license as an ID for voting [33]), etc. This makes it impossible to prevent a single person from obtaining multiple primary credentials, including ones for use by underage individuals.

Note that the alternative forms of identification suggested in [41] do not solve the problem: that paper was about accountability, not access to the Web. Here, such forms of ID could multiply the number of credentials a single person could obtain. (In 2011, President Obama's White House proposed a privacy-preserving identity scheme [40], but it said essentially nothing about how to obtain primary credentials, save that there should be many identity providers.)

By contrast, countries with a strong national ID card system are in a better position to use such schemes, especially those like Estonia which have digitally enabled ID cards [17, 11, 26]. Estonia's ID card is mandatory for all citizens

and permanent residents, and though the online version does not include birth date explicitly in a certificate field the way it should, it is part of the “personal code” field [37], which is included. Such a credential also permits a privacy-preserving registry of CL credential holders, perhaps via a Bloom filter [8], though such a database is still vulnerable to dictionary attacks by anyone who has access to the personal code. (Asserting that countries without such ID cards should adopt them is well beyond the scope of what the IETF and the W3C can accomplish. Indeed, the U.K. abandoned its recent attempt at such a system and destroyed the data [39]. In the United States, a National Research Council report described many questions that would have to be answered before the issue could even be intelligently discussed [19].)

User Challenges: Shared computers within a family pose family privacy issues. In an ideal setting, each user would have their own login and hence browser and key store. In reality, that will often not be the case. Furthermore, it may be problematic if domestic partners can see that that certificates for, say, porn sites, are already cached.

Regardless of how authentication is done, any scheme adopted must allow for use on public computers, e.g., in libraries. People do use public libraries to view age-restricted material [1], including pornography; recall, however, that social media access is also restricted in some jurisdictions. This is a thorny problem—most people do not carry devices for secure private key storage, and while online key storage systems are feasible [18], running the necessary servers is an added expense and one that is ill-borne by, say, unhoused or unemployed people who use library computers because they have no other Internet access. (While most people have smart phones, not all do, especially among the poor, the elderly, the elderly, and some minorities [31]. There are also significant issues of cross-vendor secure communication, e.g., iPhones to Windows PCs.)

Economic Issues: Operating an IDP is likely to be expensive, which raises the question of who should pay. Given that it is Web sites that must verify ages, it would seem logical that they should pay the IDPs; however, that poses its own set of problems. Naturally, Web sites will wish to minimize their own costs; however, that might impel them to select IDPs with a much narrower geographic scope, which in turn would hinder access to Web sites for many individuals. It might be possible to deal with an entirely online IDP—even in the U.S., there are online-only notary publics [28]—but that itself poses difficult governance issues, especially when crossing national boundaries.

If web sites shoulder the cost, they will have to recover it from their users. That would imply higher access charges, more ads (with their own privacy challenges), or both.

Governance: Regulating IDPs and CAs is problematic. In principle, there can be any number of either; in practice, there are issues. The first is trust: who trusts whom, to do what?

IDPs have two roles, one requiring care and judgment and the other purely ministerial. The former is at the heart of their function: they must be willing and able to judge the adequacy of the documents presented by individuals to prove their age. If we assume that there is actually harm to minors from accessing certain content—and that is the premise behind all age verification laws—then they are presumptively harmed by such accesses, and those who enabled or permitted it could be liable. By analogy to “dram shop laws,” common in at least much of the U.S. and parts of Canada, proper checking of IDs is important [23]:

When a minor is served [alcohol], you need to show that you took steps to the best of your ability to avoid purposefully serving that person. That means, in most cases, you’re only liable for selling to a minor if you didn’t ask for identification or if you looked but didn’t use your best abilities to detect the minor. If the guest presents a realistic ID that says he or she is of drinking age, no charges will be filed against the establishment.

However, if it’s obvious that the ID is fake or if that ID clearly does not belong to the person presenting it, then the establishment and the person who served that minor becomes liable.

Of course, those who wish to cheat will gravitate towards IDPs that don’t check carefully, and most people will opt for a lower-cost provider or more convenient provider, regardless of the quality of checking that is done.

The answer would appear to be regulation—but by which government? Again reasoning by analogy with alcohol, there are many governments that only loosely enforce drinking ages, if indeed they exist at all in those countries. If regulation of IDPs is by national governments, the question naturally arises as to what extent Web servers in other countries should trust them. (There is an entirely separate question, one out of scope for this essay and arguably for this workshop, of to what extent a Web server is obligated to enforce the restrictions imposed by a site visitor’s country’s laws.)

By contrast, CAs need exercise little or no judgment, save for verifying that the credentials presented are from a licensed IDP. The issue is what CAs will be accepted by Web sites, but that problem has largely been solved by the

CA/Browser Forum. This, however, is a much simpler problem than identity verification, and even web certificates have had their problems; see, e.g., [24] on Certificate Transparency as a quasi-fix. As noted, it is important that CL CAs be separate from, and not collude with, IDPs.

Regulation implies the ability for governments to audit the regulated entities' behavior. That in turn implies that logs must be kept. It is likely that such logs would include user names, addresses, ages, and forms of credentials presented. I do not explore this issue further, save to note that the properties of the CL protocol provide strong privacy guarantees against protocol-level—but not IP address-based—correlation of logs between IDPs and CAs.

Paradoxes: A security assumption behind the CL protocol is that the possessor of a subcredential's private key will not share it with others—knowledge of such a private key allows derivation of the associated primary credential's key. However, if the only use of the primary credential is obtaining age-verifying subcredentials, this isn't much of a deterrent—many people simply won't care. In other words, primary credentials need to be used for other sorts of things to create the proper disincentive for sharing.

That, however, creates pressure for mission creep. A number of other possible uses for CL credentials are given in [3], including opening bank accounts, employment verification, and vaccination certificates [29]; however, this is also a major point of social control, since it is possible to revoke a primary credential and with it all derived subcredentials [2]. Those who are disfavored by authoritarian governments may lose access not just to pornography, but to social media and all of these other services. (Of course, the need for such revocability can be a good thing, to cope with compromised credentials, lost or stolen devices, etc.)

One last note: for privacy reasons, it is important that subcredentials and client-side certificates expire relatively quickly, and that users be able to request new ones at will. It is tempting to suggest that sites that require logins for continued access only need to check a certificate once, given that age is monotonic; however, this would be susceptible to abuse by people who would set up the login with their own certificate but then give the username and password to a minor, possibly for a fee.

4 A Case Study

It is worth making some of these problems concrete. Consider a hypothetical person “Chris”, a non-driving senior citizen living with an adult child in a rural area of the U.S. To obtain a primary credential, they must first obtain an old birth certificate, possibly from out-of-state, as well as other “breeder documents” [20]. Some of these breeder documents—often, things like utility bills—may not be available to someone like Chris. Apart from the expense—quite possibly non-trivial for a poor family—Chris must persuade their child to then drive them 80 kilometers or more to a motor vehicles office to obtain a government-issued photo ID. In some states, such offices appear to be deliberately hard to access [38].

There is also the social aspect. Imagine the embarrassment to all of an older parent having to explain to their child that they wish to view pornography.

Finally, Chris probably does not own their own computer. Instead, they use their children's computer or one in a public library. Both present issues, be it privacy, credential storage, or usability.

5 Conclusion

It is worth taking an analytic look at what I've proposed here. Roughly speaking, there are three components, the technical mechanism for getting a suitable credential—here, the IDP—the mechanism for conveying it from a browser to a web site (client-side certificates), and the human and procedural aspects of obtaining an age verification credential. I examine each in turn.

CL Credentials and the IDP: Except in countries with strong, age-containing digital credentials, I assert that any solution to the age verification problem will require an IDP-like entity. That is, there needs to be some entity that will translate human-readable documents such as birth certificates or passports into a digital credential. That is part of what an IDP does. The other part, of course, is that by using the CL protocol we achieve strong privacy guarantees. That alone justifies this architecture, even if a suitable national credential is used to obtain a CL credential. Quite crucially, this privacy guarantee is a technical mechanism, not a regulatory or bureaucratic one, and hence offers much stronger guarantees of privacy.

Credential Delivery: Realistically, there are only four possible ways for a browser to deliver an age credential to a server: a new protocol to replace HTTP/HTTPS, a new HTTP header line, out-of-band communication between IDP and web sites, or in a client-side certificate. The last is clearly better. The first is absurd; the changes required by everyone are far too great. The second and third would work, but make the credential separable from the session. A Web server sees it in plaintext; a hacked web server would be fruitful source of stolen credentials. (Many of today’s age verification services essentially employ this approach. Apart from the direct privacy risk—what will such a service do with the data it collects?—there are security risks as well. For example, a database created by Tea, an app where women can warn about some bad experiences with men, was recently hacked [22]. The database contained selfies and driver’s license data; these were used to verify that the person signing up was, in fact, a woman. A corresponding app for men, TeaOnHer, has also been reported as having serious security holes [32].) By contrast, a field in a client-side certificate is strongly bound to the session and isn’t replayable without knowledge of the private key. Furthermore, such a certificate could also authenticate people after first use if desired, thereby removing friction from the user’s experience. A CA is thus desirable even without using the CL protocol, though it raises the problem of private key and certificate-sharing between devices. Possibly, the IDP-like function and the CA function could be combined into a single entity, albeit with a risk to privacy: users need a separate certificate for each site, to avoid linkage.

Human and Procedural Aspects: By far the biggest problem with the scheme I’ve outlined is the many human-scale obstacles to dealing with IDPs. *All of these apply to any other credential-based age verification scheme.* You would still need suitable breeder documents, you would still need to travel to an IDP, you would still have the multiple device and multiple forms of authentication issue, you would still have the issues of the poor, the elderly, the rural, etc., having trouble obtaining credentials. I submit that this is no different for my scheme than for *any* other, even ones that are not privacy-preserving. It is this problem above all that has to be solved, without trying to solve national-scale problems such as poverty and homelessness.

Acknowledgments

Joseph Lorenzo Hall made many useful comments on an earlier draft of this document. Merike Kaeo assisted with the discussion of the Estonian ID card. Other help was provided by Susan Landau and Ari Schwartz.

References

- [1] Steve Albrecht. “What You Need to Learn About Porn and Patron Safety”. In: *Information Today* (Sept. 2023). URL: <https://www.infotoday.com/cilmag/sep23/Albrecht--What-You-Need-to-Learn-About-Porn-and-Patron-Safety.shtml>.
- [2] Elli Androulaki, Binh Vo, and Steven M. Bellovin. *A Real-World Identity Management System with Master Secret Revocation*. Tech. rep. CUCS-008-10. Department of Computer Science, Columbia University, Apr. 2010. URL: <https://mice.cs.columbia.edu/getTechreport.php?techreportID=1421&format=pdf&>.
- [3] Elli Androulaki, Binh Vo, and Steven M. Bellovin. “Cybersecurity Through Identity Management”. In: *Engaging Data: First International Forum on the Application and Management of Personal Electronic Information*. Oct. 2009. URL: https://www.cs.columbia.edu/~smb/papers/idenman_edf09.pdf.
- [4] Noah Apthorpe, Brett Frischmann, and Yan Shvartzshnaider. *Online Age Gating: An Interdisciplinary Evaluation*. Draft. June 2025. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4937328.
- [5] Jessica Arciniega, Morgan Sexton, and Amelia Vance. *Supreme Court Upholds Age Verification: A Game-Changer for Child Online Safety Laws*. Public Interest Privacy Center. <https://publicinterestprivacy.org/paxton-age-verification/>. July 1, 2025.
- [6] *Australia Online Safety Act*. Act No. 127. 2024. URL: https://www.legislation.gov.au/C2021A00076/2024-12-11/2024-12-11/text/original/epub/OEBPS/document_1/document_1.html#_Toc185687806.

- [7] Steven M. Bellovin and Robert G. Moskowitz. *Client Certificate and Key Retrieval for IKE*. Obsolete Internet draft. Nov. 2000. URL: <https://www.cs.columbia.edu/~smb/papers/draft-ietf-ipsra-getcert-00.txt>.
- [8] B. H. Bloom. “Space/time trade-offs in hash coding with allowable errors”. In: *Communications of the ACM* 13.7 (July 1970), pp. 422–426.
- [9] Mindy Brooks. “Ensuring a safer online experience for U.S. kids and teens”. In: *Google Safety & Security blog* (July 30, 2025). URL: <https://blog.google/technology/safety-security/age-assurance-measures-safer-online-kids-teens-us/>.
- [10] J. Camenisch and A. Lysyanskaya. “An Efficient System for Non-transferable Anonymous Credentials with Optional Anonymity Revocation”. In: *Proc. of Eurocrypt ’01, LNCS 2045*. Springer-Verlag, 2001, pp. 93–118. URL: https://link.springer.com/content/pdf/10.1007/3-540-44987-6_7.pdf.
- [11] *Certificates on identity card of Republic of Estonia*. Version 3.3. Jan. 1, 2010. URL: https://www.sk.ee/upload/files/ESTEID_profiil_en-3_3.pdf.
- [12] *Crawford v. Marion County Election Board*. 533 U.S. 181. 2008. URL: <https://www.law.cornell.edu/supct/html/07-21.ZS.html>.
- [13] Roger Dingledine, Nick Mathewson, and Paul Syverson. “Tor: The Second-Generation Onion Router”. In: *Proceedings of the 13th USENIX Security Symposium*. Aug. 2004. URL: http://static.usenix.org/legacy/events/sec04/tech/full_papers/dingledine/dingledine_html/.
- [14] *EU Digital Services Act*. 2022. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32022R2065>.
- [15] *Free Speech Coalition v. Paxton*. 2025. URL: https://www.supremecourt.gov/opinions/24pdf/23-1122_3e04.pdf.
- [16] Jacob Gorman, Nikhil Mehta, Marie Nganele, Janet Zhang, and Steven M. Bellovin. *Privacy-Preserving Accountability for Non-Consensual Pornography*. In progress. 2025.
- [17] *iD*. URL: <https://www.id.ee/en/>.
- [18] Charlie Kaufman and Radia Perlman. “PDM: A New Strong Password-Based Protocol”. In: *10th USENIX Security Symposium (USENIX Security 01)*. Washington, D.C.: USENIX Association, Aug. 2001. URL: <https://www.usenix.org/conference/10th-usenix-security-symposium/pdm-new-strong-password-based-protocol>.
- [19] Stephen T. Kent and Lynette I. Millett, eds. *IDs—Not That Easy: Questions About Nationwide Identity Systems*. National Academies Press, 2002. URL: http://books.nap.edu/catalog.php?record_id=10346.
- [20] Stephen T. Kent and Lynette I. Millett, eds. *Who Goes There? Authentication Through the Lens of Privacy*. National Academies Press, 2003. URL: <http://www.nap.edu/catalog/10656.html>.
- [21] John S. Koh, Steven M. Bellovin, and Jason Nieh. “Easy Email Encryption with Easy Key Management: Why Joanie Can Encrypt”. In: *Proc. EuroSys 2019*. Dresden, DE, Mar. 2019. URL: <https://www.cs.columbia.edu/~smb/papers/eurosys-2019-submission408-e3-final-1.pdf>.
- [22] Isabella Kwai. “What to Know About the Hack at Tea, an App Where Women Share Red Flags About Men”. In: *New York Times* (July 26, 2025). URL: <https://www.nytimes.com/2025/07/26/us/tea-safety-dating-app-hack.html>.
- [23] Whitney Larson. “Serving a Minor: Who’s at Risk and What’s at Stake?” In: *Bar & Restaurant News* (Dec. 9, 2021). URL: <https://www.barandrestaurant.com/operations/serving-minor-whos-risk-and-whats-stake>.
- [24] B. Laurie, E. Messeri, and R. Stradling. *Certificate Transparency Version 2.0*. RFC 9162. IETF, Dec. 2021. DOI: 10.17487/RFC9162. URL: <http://www.rfc-editor.org/info/rfc9162>.
- [25] *Louisiana H.B. 142*. Regular session. 2022. URL: <https://legis.la.gov/legis/BillInfo.aspx?i=241701>.

- [26] *Name, Certificate, CRL and OCSP Profile for ID-1 Format Identity Documents Issued by the Republic of Estonia*. Version 1.1. Feb. 17, 2022. URL: https://www.skidsolutions.eu/upload/files/SK-CPR-ESTEID2018-EN-v1_3_20220217.pdf.
- [27] NIST. *NIST Study Evaluates Effects of Race, Age, Sex on Face Recognition Software*. Dec. 2019. URL: <https://www.nist.gov/news-events/news/2019/12/nist-study-evaluates-effects-race-age-sex-face-recognition-software>.
- [28] *Notarize Services*. URL: <https://www.notarize.com>.
- [29] Eric K. Rescorla. “A look at the Dutch vaccine passport system”. In: *Educated Guesswork blog* (Dec. 13, 2021). URL: <https://educatedguesswork.org/posts/vaccine-passport-nl/>.
- [30] Eric K. Rescorla. *The Transport Layer Security (TLS) Protocol Version 1.3*. RFC 8446. IETF, Aug. 2018. DOI: 10.17487/RFC8446. URL: <http://www.rfc-editor.org/info/rfc8446>.
- [31] Olivia Sidoti and Wyatt Dawson. *Mobile Fact Sheet*. Pew Research Center. Nov. 13, 2024. URL: <https://www.pewresearch.org/internet/fact-sheet/mobile/>.
- [32] Amanda Silberling and Zack Whittaker. “A rival Tea app for men is leaking its users’ personal data and driver’s licenses”. In: *TechCrunch* (Aug. 6, 2025). URL: <https://techcrunch.com/2025/08/06/a-rival-tea-app-for-men-is-leaking-its-users-personal-data-and-drivers-licenses/>.
- [33] Texas Department of State. *Identification Requirements for Voting*. URL: <https://www.votetexas.gov/voting/need-id.html>.
- [34] *Texas H.B. 1181*. 2023. URL: <https://legiscan.com/TX/bill/HB1181/2023>.
- [35] *UK Digital Economy Act*. c. 30. 2017. URL: <https://www.legislation.gov.uk/ukpga/2017/30/contents>.
- [36] *UK Online Safety Act*. c. 50. 2023. URL: <https://www.legislation.gov.uk/ukpga/2023/50>.
- [37] *Understand and Protect Your Digital Identity*. June 27, 2025. URL: <https://learn.e-resident.gov/ee/hc/en-gb/articles/360000633317-Understand-and-protect-your-digital-identity>.
- [38] Susan Watson. *Alabama’s DMV Shutdown Has Everything to Do With Race*. Oct. 8, 2015. URL: <https://www.aclu.org/news/voting-rights/alabamas-dmv-shutdown-has-everything-do-race>.
- [39] Brian Wheeler. “How ID Card Database Will be Destroyed”. In: *BBC News* (Nov. 10, 2020). URL: <https://www.bbc.co.uk/news/uk-politics-11719764>.
- [40] White House. *National Strategy for Trusted Identities in Cyberspace*. Apr. 15, 2011. URL: https://obamawhitehouse.archives.gov/sites/default/files/rss_viewer/NSTICstrategy_041511.pdf.
- [41] Janet Zhang and Steven M. Bellovin. “Preventing Intimate Image Abuse Via Privacy-Preserving Anonymous Credentials”. In: *SMU Science and Technology Law Review* 26 (Fall Nov. 2023), pp. 149–215. URL: <https://scholar.smu.edu/scitech/vol26/iss2/2/>.