

From Data to Duty: How Platform Knowledge Triggers Legal Responsibility

→ KEVIN FRAIZER

INTRODUCTION

Legislators at the state and federal level are exploring whether to impose age verification requirements on digital service providers, defined broadly here to refer to AI companion chatbot operators, social media platforms, and related entities. According to [New America](#), just about every state has enacted some form of an age verification law or is actively considering doing so. Members of Congress have [likewise introduced](#) myriad measures. The precise details of these proposals often are omitted from common discussion and yet merit close scrutiny. First and foremost, these policies raise important constitutional questions. Second and relatedly, the contours of age verification policies determine the extent to which users, including minors and their parents, may be required to turn over particularly sensitive information and endure time-intensive procedural processes to appeal inaccurate decisions. This brief seeks to cover the latter inquiry with a specific focus on the knowledge standard imposed on digital service providers and the corresponding implications for users.

A few qualifications are necessary in any conversation regarding technological tools and age verification. Policymaking in this field is [replete with instances of lawmakers incorrectly forecasting](#) the pace of technological advances in age verification and, critically, the extent to which those advances will maintain and enhance user privacy. [Not that long ago](#), the Third Circuit took issue with an age verification requirement given that it could not find “evidence of age verification services of products available on the market that actually reliably establish or verify the age of Internet users.” Common tools today, such as Yoti, suffer from similar shortcomings. [According to the Economist](#), “Yoti can guess the age of a white teenage boy to within less than ten months, but when faced with a dark-skinned girl is typically out by a year and a half.” Yoti’s [internal reports](#) suggest an accuracy rate of 1.1 years for users between 13 and 17. Such an error rate would produce a high number of false positives and negatives if 13 were set as a statutory cutoff for access to certain content.

Some advocates today are again counting on new technologies mitigating some of the most pressing concerns raised by families about mandatory age verification-related data disclosures. The promise of zero knowledge proofs (ZKPs) has animated a number of governments to [aggressively push for digital IDs](#). In theory, ZKPs “provide a cryptographic way to not give something away, like your exact date of birth and age from your ID, instead offering a “yes-or-no” claim (like above or below 18) to a verifier requiring a legal age threshold.”

Yet, like the hoped-for advances of previous technological eras, there’s a gulf between expectations of new privacy-protecting age-verification tools and the net result of their implementation. Somewhat paradoxically the success of ZKP technology may be what causes them to increase privacy risks. The digital service providers that deploy ZPK-based tools may collect more than simply age-based information. If Americans are forced to regularly share such information, then the aggregate impact may be to subject them to heightened risk of privacy invasions as well as increased friction in simply using the Internet. There’s also the risk that Americans may let their privacy guards down based on the supposed security of ZPK-based tools, exacerbating a tendency to share more information than necessary and to unintentionally expose oneself to privacy risks. More generally, variability in the accuracy and effectiveness of age detection technology warrants legislative humility. [Each tool](#) poses “a tradeoff between reliable accuracy and privacy.”

The legality of various age verification strategies and the technological tools therein is highly contested and unresolved. Despite the Children’s Online Privacy Protection Act (COPPA) going into effect nearly three decades

From Data to Duty: How Platform Knowledge Triggers Legal Responsibility

ago and myriad state and federal laws related to access to social media and adult content more generally being enacted in the interim, regulators and jurists have yet to adopt a consistent and precise framework of enforcement and interpretation.

The Supreme Court's recent decision in [*Moody v. NetChoice*](#) did little to clarify this ambiguity. That case dealt with facial challenges to the constitutionality of content moderation laws passed by the Texas and Florida legislatures. The Court concluded that absent far more information about how the laws may alter a wide set of possibly implicated digital service providers it could not reach a reasoned determination. Several concurring opinions left hints as to which factors may be especially important to resolving as applied challenges that may eventually come about. Justice Barrett, for instance, questioned if heightened legal protections for expressive activities may exist along a sliding scale when it comes to managing and designing a social media platform and its content feed. The upshot seems to be that the details and practical outcomes of a law may be especially important to the justices in tech cases that necessarily intersect with the First Amendment.

A subsequent decision in [*FSC v. Paxton*](#) confirmed the important role of history and practical effects on the Supreme Court's approach to laws directing digital service providers to alter the availability of certain content to certain users. The Court upheld a Texas law that mandated pornographic sites as well as other sites containing a high amount of obscene content, [*perhaps even*](#) sexual health sites, to verify the ages of users and deny access to minors. After concluding that the age verification process only incidentally burdened the ability of adults to access such content, the Court reviewed the law under intermediate scrutiny. Whereas laws that are content-based or otherwise infringe fundamental free speech rights are evaluated under strict scrutiny—meaning that they must advance a compelling governmental interest, be narrowly tailored to achieve it, and represent the least restrictive means of advancing it, intermediate scrutiny lowers the threshold for a law to be upheld. More specifically, laws that advance an important government interest and are substantially related to that objective will survive that test.

The *Paxton* Court [*slotted the law*](#) into the intermediate scrutiny framework based on several case-specific factors. Justice Thomas, writing for the majority, emphasized that “[h]istory, tradition, and precedent establish that sexual content that is obscene to minors” and that “[t]he First Amendment leaves undisturbed States’ traditional power to prevent minors from accessing speech that is obscene from their perspective.” Put differently, the law’s explicit focus on content of a certain sort—content that had long been regulated by states and subject to in-person age assurance mechanisms (namely, looking at someone’s ID)—carried considerable influence on the Court’s thinking. “No person—adult or child—has a First Amendment right to access *such speech* without first submitting proof of age,” wrote Justice Thomas (emphasis added).

Following *Paxton*, courts and scholars alike have debated the narrowness of its holdings. Some read the opinion as pertaining only to laws addressing obscene, adult content. Others think it may reach related content areas as well as design choices that may similarly conflict with commonly held views on what information is proper for children. In any event, it is highly likely that history and first principles of constitutional law and a republican form of government will sway at least a few members of the Court in related cases.

A final note before diving into the weeds of different knowledge standards: though the technological, cultural, and economic factors related to raising children to become productive members of a democratic society have shifted, a few principles have remained at the core of legislative action, judicial review, and regulatory enforcement.

From Data to Duty: How Platform Knowledge Triggers Legal Responsibility

First, parents exercise a wide latitude of authority in determining how best to raise their children and have a constitutional right to do so without direction from the state. A perceived failure by parents to make use of less invasive means than government intervention to shield their children from potentially harmful activities does not justify the latter.

Second, as stated by Senator Richard Bryan, co-sponsor of COPPA, such laws ought to “enhance parental involvement in a child’s online activities in order to protect the privacy of children in the online environment.”

Third, each technology merits unique treatment. The characteristics of the Internet and, later, social media plans that drove legislative responses are unique from those of AI. As noted above in the brief review of *Moody*, judges are particularly attuned to the idiosyncrasies of new technologies and their varied effects on foundational rights. When it comes to AI and child protection, it is clear that the unique pros and cons of AI are not widely understood. Most parents have yet to intentionally and diligently pursue greater understanding of AI. Fewer than one in five parents agreed with the following statement: “I actively seek out information and resources to better understand AI technologies.” There’s also compelling evidence that many parents have yet to meaningfully test AI tools in their personal or professional capacities. An added wrinkle in the AI space is that the relevant internal policies of AI labs are changing quickly as are the capabilities (and limitations on a per user basis) of their models. The net result is that any action taken in the AI regulatory space is likely to be out of date on several dimensions in a period of months, if not weeks.

KNOWLEDGE STANDARDS

States and federal legislators have been experimenting with different knowledge standards related to a user’s age to trigger additional legal responsibilities for digital service providers. A full analysis of those legislative interventions is beyond the scope of this overview. The more generic analysis adopted here suffices to surface the most important takeaway: selection of a knowledge standard has significant and unavoidable consequences in altering the likely behavior of users, parents, and digital service providers.

ACTUAL KNOWLEDGE

A digital service provider has actual knowledge of a user’s age when they are directly informed or consciously aware of that status. A dictionary definition of the term which some courts have relied on is “direct and clear knowledge.”

Yet, for reasons detailed below, this standard is commonly supplemented with a prohibition on digital service providers intentionally avoiding learning a user’s age so as to sidestep any requirements triggered by the actual knowledge standard. Some definitions of actual knowledge explicitly prohibit a provider from willfully disregarding a user’s age. Pursuant to such definitions, “persons who know enough to blind themselves to direct proof of critical facts in effect have actual knowledge of those facts.” Under other statutes, such as in the Colorado Privacy Act, willful disregard is not implied by the actual knowledge requirement but rather listed as an additional, separate standard. Under the Florida Digital Bill of Rights, willful disregard is tantamount to actual knowledge. In practice, foreclosure of willfully disregarding a user’s age means that if a provider has “access to information that could inform them that a user is under 18 years-old and chooses not to act on it, they may still be held responsible” under the applicable law. Courts have often turned to a two-part test for this inquiry: “(1) the defendant must subjectively believe that there is a high probability that a fact exists and (2) the defendant must take deliberate actions to avoid learning of that fact.”

From Data to Duty: How Platform Knowledge Triggers Legal Responsibility

Actual knowledge is a common knowledge standard across federal and state regulations and proposals. [Several aspects of COPPA](#), for instance, are triggered upon a digital service provider having actual knowledge that a user is under the age of 13. Various states have integrated this standard into relevant laws including but not limited to [Colorado](#) and [Arkansas](#), [although the latter was enjoined](#) by a federal district court. One analyst estimates that of comprehensive state privacy laws with youth privacy provisions [the vast majority](#) (75 percent) rely on an actual knowledge standard. Many state laws impose various age related provisions yet omit an explicit knowledge standard. In those instances, [it's likely that](#) an actual knowledge standard will apply.

From the perspective of a digital service provider, this standard presents somewhat conflicting incentives absent some clarifying language as to how regulators may treat willful disregard of a user's age. Providers may intentionally avoid learning any information about users' ages so as to avoid having such knowledge. This cuts against the legislative spirit presumably animating age verification laws. Those providers that adopt a more proactive stance risk being found non-compliant if they do not implement the necessary safeguards upon determining a user is a minor.

From a user's perspective—assuming, [as is the case with COPPA](#), that providers may not circumvent an actual knowledge standard ignoring a user's age—this standard will likely subject them to comparatively few disclosure requirements relative to other standards, as covered below. Under an actual knowledge framework, providers generally are not required to deploy age verification tools across their entire user base. Instead, duties attach when the provider receives clear notice of a user's age, whether through user self-identification, parental contact, or internal reporting mechanisms. Because there is no categorical obligation to screen every user, the vast majority of users will not be asked to provide government identification, biometric data, or additional attestations. In this respect, the standard minimizes friction and data collection relative to constructive knowledge or “reasonably should know” approaches that encourage broad-based monitoring.

CONSTRUCTIVE KNOWLEDGE

A constructive knowledge standard expands the inquiry beyond those instances in which a digital service provider is directly on notice of a user's age and [instead asks](#) “what an operator ought to know about its users if they have carried their work in due diligence.” The focus is objective and grounded in the conduct expected of a reasonable operator under the circumstances. The analysis centers on whether a provider, exercising ordinary diligence and drawing on the data it collects and uses in the normal course of business, would recognize that a user is a minor. This formulation shifts the compliance inquiry from isolated moments of express disclosure to the design of systems, internal review processes, and the treatment of user signals. As a result, liability may attach based on how a platform structures onboarding flows, configures content recommendations, or responds to age-related indicators arising from user activity or platform operations, such as advertising.

Constructive knowledge formulations have appeared in both enacted statutes and high-profile legislative proposals, often through “knew or should have known” language or standards that hinge on what a reasonable operator would identify through ordinary business practices. California's [Age-Appropriate Design Code Act](#), for instance, applies to online services “likely to be accessed by children,” a formulation that does not require direct notice of a particular user's age. Similar approaches appear in other state youth privacy and social media measures that tie compliance obligations to audience composition, platform design features, or patterns of use that suggest a service attracts minors.

At the federal level, several bills introduced in recent Congresses have [incorporated variations of this constructive framing](#), either by imposing duties when a provider “should reasonably know” that minors are

From Data to Duty: How Platform Knowledge Triggers Legal Responsibility

using the service or by defining covered services based on their expected child user base rather than confirmed individual ages. These formulations reflect a legislative judgment that reliance on self-reporting mechanisms for users to state their age may be insufficient to advance child-protection goals.

Litigation involving these constructive-style triggers remains ongoing and fact-dependent. [Challenges to the California Age-Appropriate Design Code Act](#), for example, have focused in part on whether obligations tied to likely child access effectively compel broad age screening or redesign of core product features. These kinds of regulatory spillovers implicate core privacy and expression interests. Courts evaluating such provisions have signaled concern where compliance pressures push platforms toward universal age verification or where statutory language lacks precision as to what diligence is required. The Due Process Clause of the Fourteenth Amendment [shields individuals and entities from vague laws](#) in which the government has [excessively broad enforcement discretion](#) and the public has insufficient guidance as to how to avoid prosecution, penalty, or both.

The broader takeaway from these early disputes is that constructive knowledge standards expand both the scope of covered services and the evidentiary questions courts must resolve. Rather than asking whether a provider was directly informed of a user's age, courts must assess what signals were available, how platforms processed those signals, and whether the law, in operation, incentivizes pervasive data collection. In that respect, the doctrinal and practical consequences of constructive knowledge are considerably more complex than those associated with an actual knowledge trigger.

From the perspective of users, a constructive knowledge regime alters the background conditions of participation online. When compliance turns on what a provider ought to know through reasonable diligence, providers have incentives to monitor age-related signals more closely and to document how they respond to them. [That shift can produce](#) a subtle but [meaningful chilling effect](#). Users, including adults, may [become more cautious](#) about posting content, joining communities, or visiting certain sites if they understand that their activity could be interpreted as evidence of age or as triggering additional scrutiny. Minors, in particular, may [avoid searching for sensitive information](#) related to health, sexuality, or identity if they believe those queries will be flagged as indicators of age and lead to restrictions or reporting. Even adults may hesitate to engage with lawful but controversial content if doing so increases the likelihood of verification prompts or account limitations. This may even cause adults to forgo pro-social actions, such as watching kid-related content to see if it aligns with their family's values, because doing so may sort them into a "minor" distinction--thereby requiring them to prove they're indeed an adult via some form of age verification.

The record surrounding age assurance laws underscores these risks. Advocacy organizations have warned that broad age-verification or age-estimation requirements can deter individuals from accessing lawful speech because they do not wish to submit identifying information or risk erroneous classification. The same dynamic applies in a constructive knowledge framework that encourages providers to infer age from behavioral data rather than from explicit self-disclosure. When the cost of error includes account suspension, content removal, or the need to appeal a misclassification, rational users may self-censor. These practical effects are difficult to quantify, but they are central to the constitutional analysis previewed above. As courts have emphasized in related contexts, the real-world burdens imposed by a regulatory scheme, including the friction and stigma associated with verification, bear directly on whether the law unduly restricts protected activity.

REASONABLY DETERMINE

A "reasonably determine" standard requires a digital service provider to reach an age-related conclusion based on objective factors and reasonable operational steps, even in the absence of direct notice. Under this approach,

From Data to Duty: How Platform Knowledge Triggers Legal Responsibility

the obligation is triggered when a provider has reasonable grounds to conclude that a user is below the relevant age threshold after applying defined criteria to the information available to it. Variants of this formulation appear in statutes using phrases such as “reasonably believes,” “reasonable grounds to believe,” or “commercially reasonable efforts” to ascertain age. In some laws, the standard is further qualified by reference to the “level of certainty appropriate to the risks” associated with the service, thereby linking the rigor of age-determination efforts to the nature of the product and the harms at issue. Other formulations, including aspects of the COPPA mixed-audience framework, require operators to use means “reasonably calculated” to determine whether a visitor is a child before collecting certain information.

Although this standard shares vocabulary with constructive knowledge, it is analytically distinct. Constructive knowledge turns on what a provider ought to know if it carried out due diligence, a formulation that invites retrospective debate over whether age could have been inferred from a diffuse set of signals. A “reasonably determine” framework is more often structured as a forward-looking decision rule. It asks whether the provider applied identifiable criteria and reached a reasonable conclusion based on those criteria. Several state enactments illustrate this distinction.

Louisiana’s social media law, for example, applies when a platform “reasonably believes or has actual knowledge” that a user is under sixteen and requires “commercially reasonable efforts” calibrated to the risks posed by the service. Similarly, certain legislative proposals direct regulators to specify acceptable methods and accuracy thresholds, taking into account technical feasibility, cost, and the resources of covered entities. These provisions expand the universe of information a provider is deemed to know. They also attempt to structure how the provider makes and documents an age-related judgment. However, there are limits as to how abstract a legislature may be in directing digital service providers to assess the age of users. By way of example, in NetChoice v. Carr, a federal district court determined that a Georgia requirement that social media platforms “make commercially reasonable efforts to verify the age of account holders” was too ambiguous. It was unclear to the court whether such vague mandates would impose different requirements on different digital service providers given that what’s commercially reasonable for one firm would likely not be for another. The state admitted such an outcome was feasible given the text of the law. Note this decision is currently on appeal.

The structural features common to “reasonably determine” requirements is what makes the standard sufficiently distinct to merit separate analysis. To be sure, there is overlap in practical consequences. Both constructive knowledge and “reasonably determine” approaches may encourage providers to collect additional signals, formalize internal review processes, and document compliance decisions. Yet the legal inquiry differs. Under a constructive knowledge theory, litigation may focus on whether the provider should have inferred age from available data. Under a “reasonably determine” framework, courts are more likely to examine whether the provider employed reasonable criteria, followed the prescribed process, and acted consistently with any specified statutory factors. That difference has implications for vagueness challenges, for evidentiary burdens, and for the extent to which compliance pressure evolves into a de facto mandate to verify the age of all users.

From the perspective of digital service providers, a “reasonably determine” standard is likely to influence product design in more structured and visible ways than either actual or constructive knowledge. Because the inquiry centers on whether the provider applied identifiable criteria and followed reasonable procedures, platforms have incentives to formalize age-related decision points within onboarding flows, account settings, and content gating mechanisms. Providers may develop tiered systems in which certain features are limited until age is determined through specified methods, or in which particular risk indicators trigger escalation to additional review. This can include building internal checklists, documenting the factors considered in reaching

From Data to Duty: How Platform Knowledge Triggers Legal Responsibility

an age determination, and calibrating the degree of certainty to the nature of the content or service at issue. Where statutes reference “commercially reasonable efforts” or link the required rigor to the risks posed by the platform, providers may adjust their data collection and verification tools accordingly, developing more thorough verification flows in higher-risk contexts and adopting lighter-touch measures elsewhere. In that respect, the standard encourages proactive system design and compliance documentation, rather than the more diffuse monitoring incentives associated with constructive knowledge.

For users, the ramifications are mixed. On the one hand, a clearly structured “reasonably determine” framework can reduce some of the unpredictability associated with open-ended inference standards. If platforms rely on enumerated factors and defined escalation procedures, users may encounter more transparent age prompts and clearer explanations of why certain features are restricted. On the other hand, the very existence of formal decision criteria may increase the frequency of age-related interventions. Users could face additional prompts at sign-up, periodic re-confirmation requests, or feature-specific verification requirements tied to perceived risk. As with other knowledge standards, the risk of error remains salient. A user incorrectly classified as a minor may experience account limitations and bear the burden of appealing that decision, which may not be a straightforward process despite provider assurances. Moreover, if age determinations are informed by contextual or behavioral signals, users may adjust how they post, search, or interact in order to avoid triggering additional scrutiny. Even a well-calibrated “reasonably determine” regime thus has the potential to introduce friction and to shape user behavior in subtle ways that extend beyond the narrow goal of identifying minors.

Legislatures do not always confine themselves to the three formulations discussed above. Variants appear across state and federal proposals that borrow from familiar culpability concepts such as “reason to know,” “knew or should have known,” or “reckless disregard.” Others blend scope-based triggers, such as whether a service is “likely to be accessed” by minors, with user-specific knowledge standards. These formulations can impose distinct obligations. A negligence-style “should have known” provision may expand exposure by inviting inquiry into what signals were available but not acted upon. A recklessness or willful blindness standard may focus on conscious disregard of age-related risks and increase the importance of internal documentation. Scope-based triggers may attach duties even in the absence of individualized knowledge about a particular user. Each of these approaches can meaningfully alter how providers design compliance systems and how users experience friction.

At the same time, the practical consequences of these variants often map onto the same core tensions already examined. Negligence-style language resembles constructive knowledge in its incentive to monitor and infer age from available data. Recklessness and willful blindness echo the concerns raised in the discussion of actual knowledge about strategic ignorance and documentation. Scope-based triggers raise many of the same constitutional and operational questions implicated by probability-based “reasonably determine” standards, particularly where they pressure platforms toward broad screening or redesign of core features. The doctrinal labels differ, and in litigation those differences can matter. Yet the behavioral incentives they create tend to cluster around a limited set of trade-offs.

For that reason, this brief focuses on actual knowledge, constructive knowledge, and “reasonably determine” as the principal frameworks through which legislators structure age-related obligations. Together, these standards capture the central policy choices: whether to require direct notice, to impose an objective duty to infer, or to mandate a structured determination process. They also surface the most salient tensions between child protection, parental authority, user privacy, platform design autonomy, and constitutional constraints.

From Data to Duty: How Platform Knowledge Triggers Legal Responsibility

Variants deserve attention in careful statutory drafting, but the big three provide the clearest lens for evaluating how knowledge standards reshape the online environment for both users and digital service providers.

CONCLUSION

“Minors are entitled to a significant measure of First Amendment protection, and only in relatively narrow and well-defined circumstances may government bar public dissemination of protected materials to them,” as declared by the Supreme Court in *Brown v. Entertainment Merchants Association*. The Court additionally made clear that while there is “no doubt a State possesses legitimate power to protect children from harm,” such power is not unlimited. Specifically, that authority “does not include a free-floating power to restrict the ideas to which children may be exposed. Speech that is neither obscene as to youths nor subject to some other legitimate proscription cannot be suppressed solely to protect the young from ideas or images that a legislative body thinks unsuitable for them.”

Given that explicit protection of the First Amendment rights of young Americans, legislators ought to thoroughly assess which of the aforementioned knowledge standards is appropriate in any age verification legislation. The burdens on legislators do not end there. Any selected knowledge standard may also implicate the First Amendment rights afforded to digital service providers. Given that the Supreme Court has yet to definitively state whether platform design decisions constitute expressive activity that warrants heightened constitutional protection, legislators ought to tread carefully when directly or indirectly pushing providers to substantively alter how, why, and to whom they present content. And, legislators must also pay close attention to the fact that each knowledge standard imposes some degree of data exchange between users and providers--and perhaps third parties.

In sum, the selection of a knowledge standard by legislators will determine how platforms structure their systems, interact with users, and rely on third parties for additional assistance, if at all. An actual knowledge regime confines obligations to situations of clear notice and limits incentives for pervasive screening. By contrast, constructive or “reasonably determine” standards pressure providers to collect additional signals, formalize inference processes, and document compliance decisions in ways that can approximate universal age surveillance. Those architectural consequences affect not only minors but all users. Legislators should therefore evaluate knowledge standards not only by reference to abstract child-protection goals, but also by their predictable effects on data collection, anonymity, and access to lawful speech.

The Court’s decision in *FSC v. Paxton* does not resolve these broader questions. That case addressed content historically treated as obscene as to minors and assessed a specific verification regime in that context. It does not supply a general license to impose age verification across social media platforms, AI systems, or mixed-content services where minors retain full First Amendment protection. Age assurance regimes operate through third-party vendors, data flows, and technical systems that introduce new privacy and cybersecurity risks. A statute may appear narrowly tailored on its face yet produce diffuse and lasting burdens in practice. Careful attention to those downstream effects is essential if child-protection efforts are to remain consistent with constitutional limits and sound governance.

**GET IN
TOUCH**

🌐 abundance.institute
✉️ [@abundanceinst](https://twitter.com/abundanceinst)
📧 @abundanceinstitute

Kevin Fraizer
Senior Fellow
kevin@abundance.institute