Online Age-Verification: Protecting Children and Privacy

Adam Candeub, Senior Fellow
Center for Renewing America

The State of Utah passed landmark legislation with its Utah Social Media Regulation Act, S.B. 152 (the "Act"), which returns to parents their traditional authority over their children's education and upbringing by requiring parental permission for minors to open social media accounts. The States of Arkansas, Louisiana, Texas, and Virginia also subsequently passed similar laws.[1]

Some of these laws specify acceptable methods of verifying age. Some, like Utah's, delegate this task to administrative agencies. This white paper examines the legal requirements that authentication methods should satisfy and current verification methods and technologies. In particular, the white paper examines technologies that exploit the features of zero-knowledge proofs. These technologies allow internet users to verify a feature about themselves, such as age, citizenship, or insurance status, without revealing anything else, therefore remaining anonymous.

State authentication regimes will have consequences beyond age verification. The next big thing in internet security will be proof of humanity. As if *Blade Runner* is coming to life, internet users must prove they are people, not AI. This reality exists as "complete this puzzle to prove you are not a robot." States now have the opportunity to become the first to create a best-in-class humanity verification regime as an outgrowth of its age verification regime that preserves privacy, respects our values, and promotes the trust needed for a thriving internet ecosystem.

The paper makes four points:

- Legal challengers to age verification laws will likely look to *Ashcroft v. ACLU,* in which the Supreme Court said the Child Online Protection Act (COPA) was unconstitutional, finding that (1) age verification for access to pornographic websites burdens adult speech under the First Amendment and (2) filters are more effective than age verification. However, because *Ashcroft* involved internet access, not contract or account formation, as do current state age verification laws, it does not control. Further, it stands on certain factual predicates concerning the effectiveness of filters that two decades of online experience debunk. Nevertheless, out of an abundance of legal caution, the states can minimize the risk of overburdening adult speech by allowing numerous age verification methods that let individuals and online firms choose which work best for them.

- States should recommend a non-exclusive list of verification techniques. In the absence of perfect functionality, a reasonable regulatory goal is to raise the cost of circumventing verification with the expectation that increasing the cost of undesirable behavior will decrease its prevalence.
- States, in setting a menu of acceptable methods, should encourage firms to develop age verification techniques that employ zero-knowledge proofs (ZKPs), a mathematical concept several decades old but only recently applied to online settings, particularly cryptocurrency. As we explain, these and similar techniques, such as digital signatures, offer a high degree of privacy at minimal cost.
- Age verification is but a part of the growing need to authenticate human beings on the internet in the face of ever-growing AI capabilities. Technologies that employ ZKPs and similar techniques, such as digital signatures, offer a privacy-respecting answer to this challenge. States can become leaders by encouraging the development of authentication regimes that rely on ZKP technologies.

Protecting Parents' Right to Supervise Their Children's Upbringing While Minimizing Burdens on Adult Speech

Western societies have long understood how literature, music, images, and other cultural materials influence children. Plato asserted in *The Republic* that music composed in the mixed Lydian mode results in moral weakness, while the Phrygian mode furthers courage among young men. Mirroring Plato's concern, today's parents and teachers devote tremendous time and energy to selecting books for school curricula and libraries. Given that books and other media can have either a positive or negative effect, parents have always worked to ensure that children should be exposed to salubrious cultural materials and avoid the hurtful or distracting.

The harms of social media are not strange or improbable as modern readers might consider Plato's judgment about the mixed Lydian mode.[2] Depression, self-harm, suicide attempts, and suicide all increased sharply among U.S. adolescents between 2011 and 2019,[3] with similar trends worldwide.[4] The increase occurred at the same time that use of social media skyrocketed, becoming a fixed, essential feature in teens' lives.[5] Social media is a prime suspect for the sudden rise in mental health issues among teens.[6] This suspicion has been borne out by studies positing a causal role between social media use and decreased emotional well-being.[7]

The Supreme Court has recognized the importance of respecting and supporting parental authority over what their children see and hear. Parents have the right to educate and protect their children from unwanted cultural influences, a right the Court recognized in *Pierce v. Society of Sisters* as central to a democratic society. "The fundamental theory of liberty upon which all governments in this Union repose . . . [is that] the child is not the mere creature of the State; those who nurture him and direct his destiny have the right, coupled with the high duty, to recognize and prepare him for additional obligations."[8] The Court later wrote that in "light of [its] extensive precedent, it cannot now be doubted that the Due Process Clause of the Fourteenth Amendment protects the fundamental right of parents to make decisions concerning the care, custody, and control of their children."[9]

The Court has upheld laws that protect this parental authority in the face of ever more intrusive technologies that make it easier to reproduce and distribute text and images. For instance, in *FCC v. Pacifica*, the Court ruled that the Federal Communications Commission may regulate the content of radio broadcasts to ensure that children are not exposed to indecent content.[10] The FCC's indecency regulations are still in effect today.[11]

Similarly, the Supreme Court recognized in *Rowan v. U.S. Post Off. Dept.*[12] that parents have the right to block unwanted solicitations and communications that could harm children. There is no First Amendment right to communicate with children against their parent's wishes. If parents can prevent a mailer from sending paper solicitations to their kids, consistent with the First Amendment, the state may limit online analogs.

On the other hand, the Court in 2004 struck down the Child Online Protection Act in the landmark case *Ashcroft v. ACLU*,[13] which required age verification such as a credit card to view online pornography. Because the law was content-based, the Court applied First Amendment strict scrutiny, which requires Congress to identify a compelling government interest and narrowly tailor its prohibitions to address that interest without affecting others, like protected speech. The Court found that "filters are more effective than age-verification requirements," therefore, age-verification is overly broad and burdensome.[14] But, the manifest ineffectiveness of filters and the introduction of smartphones, which became widespread a decade after the ruling, obviate this ruling's factual predicates and bring its precedential value into question.

Moreover, social media age verification laws, unlike COPA, do not authenticate for the purpose of age verification to view pornography—or any other content. Instead, it requires age verification for a platform to enter into a contractual relationship when forming an account. It is a content-neutral regulation of contract law.

As a general rule, all contracts by a minor are valid but voidable with certain exceptions.[15] And even though a minor can void most contracts he enters into, most jurisdictions have laws that hold a minor accountable for the benefits he received under the contract.[16] Because children can make enforceable contracts for which parents could end up bearing responsibility, it is a reasonable regulation that parental consent would be required for such contracts. While few courts have addressed the question of the enforceability of online contracts with minors, the handful of courts that have, held the contracts enforceable on the receipt of the mildest benefit.[17] Due to the inevitable parental interest in the contracts into which their children enter, both the Constitution and state laws require parental consent for all sorts of contracts and agreements--ranging from getting a tattoo,[18] obtaining healthcare,[19] marriage,[20] obtaining a driver's license,[21] entrance into the military service,[22] waiving the right to counsel,[23] or using a tanning facility.[24]

At the same time, even content-neutral laws, such as the Utah Social Media Regulation Act, S.B. 152, might be held, by an ideologically motivated court, unconstitutional if it incidentally inhibits or burdens constitutionally protected speech.[25] And, the Court in *Ashcroft* ruled that age-verification systems burden speech because they "require all "adults . . .[to] gain access to speech they have a right to see . .. having to identify themselves or provide their credit card

information."[26] Thus, in an abundance of legal caution, states should place as light a burden as possible on adult internet users.

Because courts might view age verification as a burden on speech, the Department should specify numerous types of age verification and accept unenumerated methods that function equivalently. This will allow firms and individuals maximum choice in finding verification that works for them—without overburdening the speech of adults.

From time immemorial, kids have gotten around age verification legal requirements. Fake ids, cooperating older siblings or friends, and other techniques have typically allowed minors to circumvent age restrictions, from getting into nightclubs and bars to buying alcohol and cigarettes or Playboy magazines at the drugstore. State laws or regulations cannot achieve absolute prohibition without astronomical costs and imposition on civil liberties. Intelligent regulation will then raise the cost and increase the difficulty of unlawful underage access, hoping to shift overall behavior significantly and positively.

Age Verification Methods

There is a wide variety of potential age verification methods. Regulations that allow the greatest number of options and give individuals and firms choice have the best chance of surviving legal scrutiny. The less burdensome age verification is for adults, the greater chance of surviving legal challenges. The following are various possible methods as well as a generic description of *any* acceptable method.[27]

Providing identity documents

The most straightforward age verification method requires furnishing of government-issued identification or financial document. Users could furnish pictures of driver's licenses or similar documents or financial records. This is hardly a radical or unusual requirement as it already applies to online gambling and dating apps and online ordering of cannabis products, cigarettes, and alcohol.

Some other jurisdictions are applying age verification to online access. For example, the UK's Online Safety Act,[28] expected to become law this fall, requires, as one method to verify age, uploading the details of a form of identification.[29]

It should be noted that social media companies need not be the entities that provide the verification. Germany uses this approach. According to German law regarding online access for minors, the State Treaty on the Protection of Young Persons in Broadcasting and Telemedia (the JMStV), specific content, which is harmful to minors, may only be distributed through the internet if the distributor ensures that only adults have access through closed age verified user groups ("AV systems"). According to press reports, the German Commission for the Protection of Minors (KJM) has certified 99 concepts and modules for AV systems. For instance, Incode Technologies' and Veriff's solution verifies the users' age by validating their ID card against a selfie and/or live video.

Beyond government-issued IDs, financial documents could be used to verify age. Indeed, that was the requirement used in COPA, the law the Supreme Court declared unconstitutional in *Ashcroft*.[30] A credit card number, debit card, or bank account number could serve as age verification, as underaged individuals generally cannot have credit cards, open bank accounts, or obtain similar financial services. If an individual pays by a credit card, an online service can reasonably assume that the credit card company effectively verified the consumer, and thus the purchase is allowed.

Biometric Age Markers

Technologies can use AI analyses of images of our faces or retinal scans[31] to estimate age. Companies such as FaceTec and Yoti are already developing the commercialization of these techniques. While these techniques can be extremely accurate, especially when used in conjunction with government IDs that have pictures, they pose significant privacy concerns. Even if all biometric data were analyzed and retained by a third party, many people would balk at regularly (or even irregularly) providing photographs or retinal scans online.

Behavioral approaches

"Behavioral" approaches allow a firm to assess the user's internet use patterns and, using data analysis or AI, make an educated guess as to the user's age. This approach requires no sharing of personal information (e.g., a user's actual date of birth, driver's license number, or other data) or documents. While this method does require the collection and analysis of tracking data, most internet platforms already do this. This approach would burden those who use the internet anonymously through browsers like TOR or select no tracking privacy settings. The UK is implementing "estimation measures," i.e., what appear to be behavioral approaches, to prevent children from accessing pornography in implementing its soon-to-be-passed Online Safety Act.[32]

Technologies Implementing Zero-Knowledge Proofs and Digital Signatures

One of the most interesting technologies that firms could use to comply with age verification requirements employ zero-knowledge proofs (ZKP). Previously more common in mathematics, ZKP technology and protocols are beginning to be implemented in a wide variety of online and internet scenarios.[33] For example, ZKP applications currently exist and are being used in blockchain and cryptocurrency.[34] Indeed, in a simple and analogous way, public-private key privacy relies on some of the mathematical features they employ.

The concept they employ is simple to explain by intuitive metaphor, while the mathematics is a bit more complicated.

Metaphoric Introductions to Zero-Knowledge Proofs

The basic idea behind zero-knowledge proofs is that a "prover" demonstrates to the "verifier" that he knows some fact. In practice, the "fact" could be age, citizenship, or insurance status and is encoded in a number that solves some mathematical operation. If the user knows the number,

he can demonstrate the fact. Significantly, he can show he knows the number (or fact) without revealing it. That is a somewhat odd concept, but the following two examples illustrate it metaphorically and non-mathematically.

Ali Baba's Cave

This is one of the earliest examples put forth by one of the pioneers in the area, the Belgian cryptographer Jean-Jacques Quisquater.[35] Imagine a U-shaped cave with a north and south entrance. At the back of the cave, there is a door with a padlock. The Prover must show the Verifier that he knows the padlock's combination without revealing the actual combination.

To do so, the Verifier shouts to the Prover (who is in the cave), "Appear in the North entrance." There are three possibilities: the Prover is (luckily) on the north side of the cave, in which case he appears. Or, the Prover is on the south side, knows the combination, passes through the door, unlocks the padlock, and appears on the north side. Last, the Prover doesn't know the combination and stays on the south side of the cave.

In that last instance, he cannot appear on the North side. He is revealed as fraudulently claiming to know the combination.

But, if the Prover appears on the north side, then the Verifier knows that he is either (i) lucky because he was already there or (ii) was on the south side of the cave and could appear in the north entrance because he knows the combination and could pass through the door. The Prover then says "Appear at the South entrance." If the Prover does so, the Verifier can be confident the Prover knows the combination—without ever learning it.

Proving Balls are Differently Colored to a Color-Blind Verifier

Your friend, the Verifier, is colorblind. There are two balls; one is red, the other is blue. You (the Prover) want to prove to your colorblind friend, the Verifier, that the balls are different colors without identifying one as red and the other as blue.

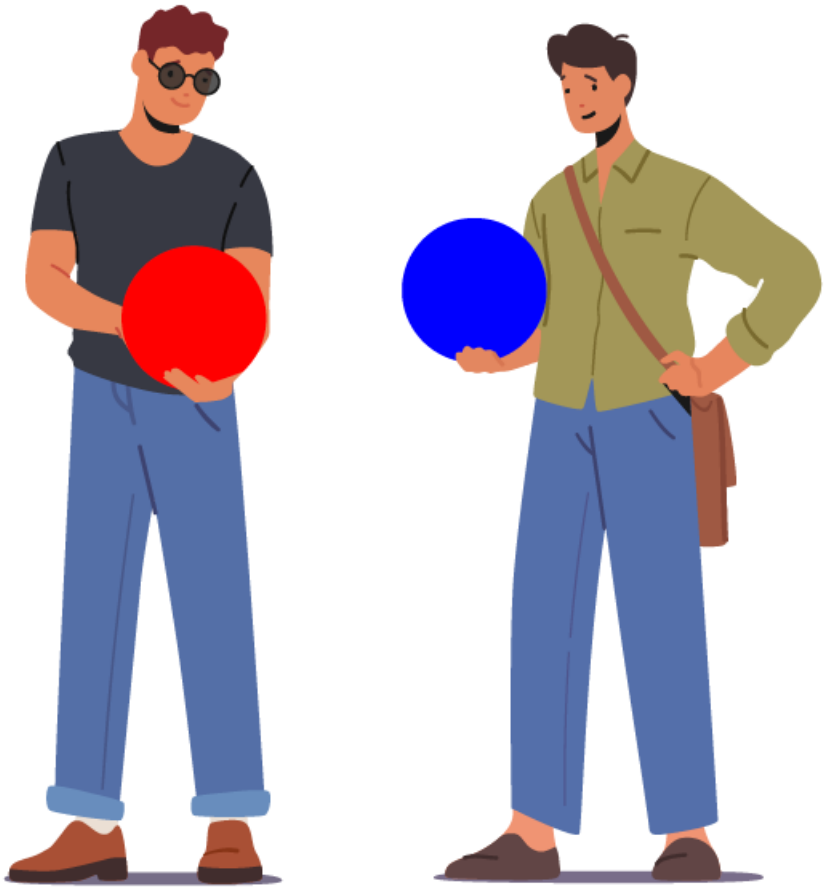

How can the Prover show his colorblind friend, the Verifier, that the balls are different colors? Easy. The Verifier places the balls before the Prover in a specific order. Red Ball is on the left; blue ball is on the right. He then picks up the balls from the table and places the balls behind his back so the Prover cannot see.

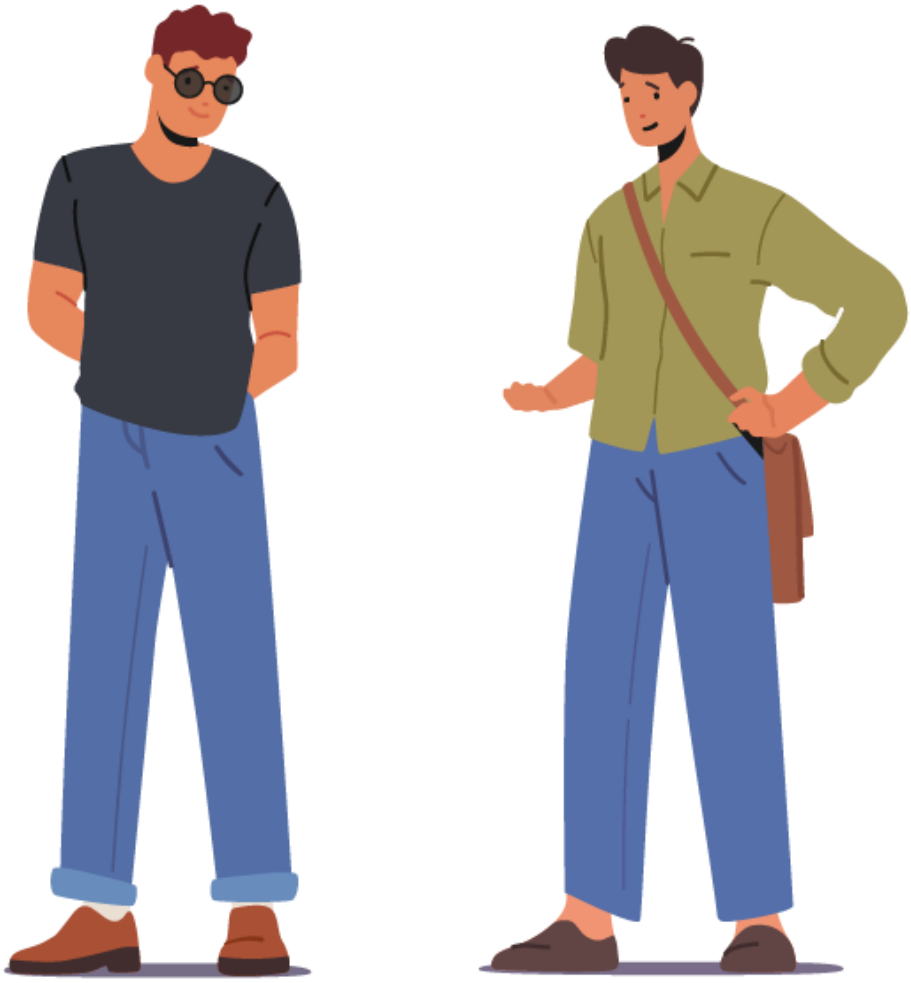

The Verifier then returns the balls to the original placement and gives the Prover a challenge: he asks the Prover if he switched the places of the balls.

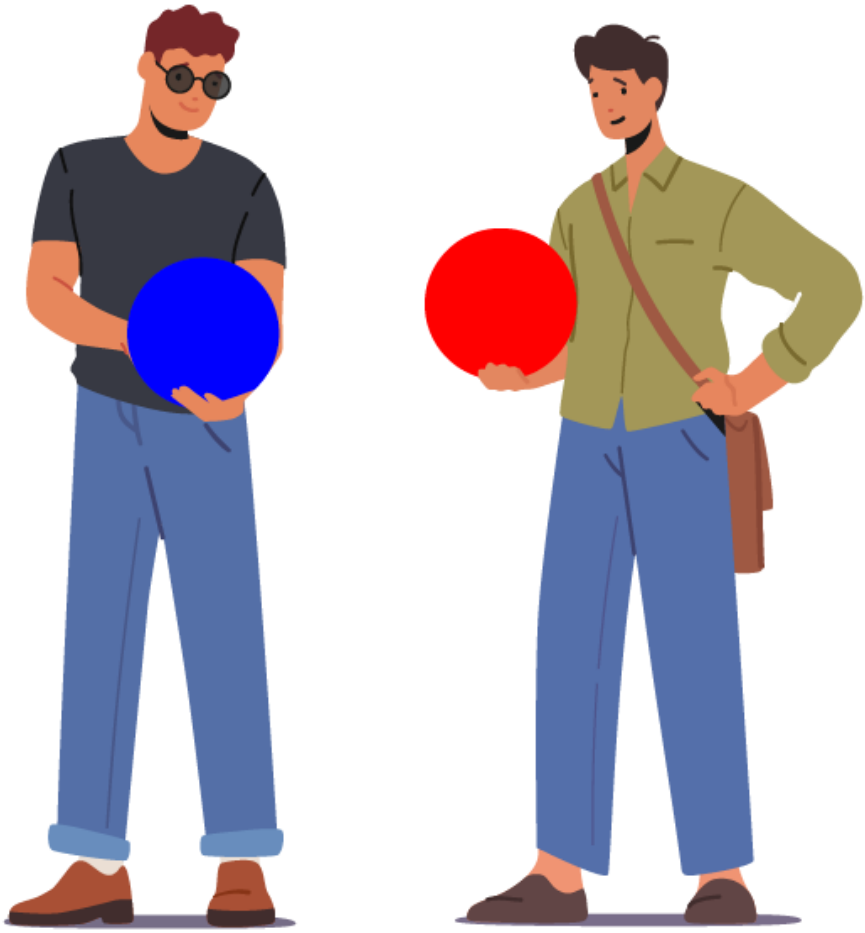

If the balls are the same color, then the Prover will never say that the Verifier switched. If they are different colors, the Prover will be able to tell the Verifier when he switched them behind his back—and when he did not.

Of course, the Prover could guess whether the Verifier switched and would be right ½ of the time. However, this strategy could not be repeated with success within normal probability ranges. Thus, the Prover can show the Verifier that he knows the color of the balls—without revealing which ball is colored blue or red.

Layman's Introduction to Public-Private Keys, Digital Signatures, and Zero-Knowledge Proofs

Rather than show that he knows some fact, i.e., that the balls are different colors, ZKP can be used to show that a Prover knows virtually any fact. At a high level, this is the way it would

work: A trusted third party could provide a user who wishes to set up an account (the Prover) with a secret number. The number could represent anything, i.e., citizenship, credit score, health insurance status, car registration status—or of particular relevance here, age. These could be established through any of the means described above.

The trusted third party could share with the platform with whom the Prover wants to set up an account (the Verifier) a challenge, which is a mathematical problem that can only be solved (with currently known computation speed) with the secret number. If the user passes the challenge, she is authenticated as a certain age, citizenship, credit score, etc. The point is that the user authenticated herself without ever revealing to the Platform anything about who she, in fact, is--and didn't even reveal the secret number that could connect her to her identity via the trusted third party. The user has complete anonymity with regard to the Platform.

Obviously, anonymous online age verification could use this strategy. If an online person (the "Prover") can prove to a platform that he is above the required age without revealing anything else about himself, that would allow age verification without compromising privacy. The following attempts to give a brief introduction into how computers can do this trick with numbers. It requires a little background through public-private key encryption and digital signatures.

Public Keys and Symmetric Encryption

When most people think of encryption, they think of symmetric encryption. For example, in symmetrical encryption, like the so-called "Caesar cipher supposedly used by Julius Caesar, you use a function: replace a letter in the message to be encrypted with a letter a fixed number of positions down the alphabet. For example, if the fixed number were "3," "AVE" would be encrypted in the following way: A would be replaced by D, V would become Y, and E would be H— "DYH." In order to decode the message, Caesar would have to send you the key, "3."

The Caesar cipher is additive. The fixed number a/k/a/ "key" you used to encrypt was "+3" and to decrypt was -3. Thus, both the encrypting and decrypting keys are known to both the sender and receiver; one is just the additive inverse. They are both "public keys."

But, let's say the function was multiplicative. To encrypt, you use your fixed number X in a function or algorithm, and to decrypt, you use X's multiplicative inverse, 1/x. And, let's say you could select x and 1/x so that it is impossible to derive x from 1/x. (More on how that is done in the next subsection.) You can share one key, making it public, and keep the other private. If X is your public key and 1/x your private key, anyone can encrypt a document using X—but only you can decode using 1/x. And you can encrypt a document with your private key, and anyone can decrypt it with your public key. While such keys would be multiplicative inverses, they of course would *not* be x and 1/x, because anyone could derive one from the other. How to find numbers that are multiplicative inverses of each other such that it is easy to invert one way but hard to un-invert, i.e., run the function the other way, is an impressive mathematical feat discussed below.

Math Problems that are Hard in Only One Direction and Public-Private Key Encryption

Consider this problem. Suppose you only had 10 minutes and a pad and pencil to get the result: x $=13^9$. You probably could figure out x (10,604,499,373) within ten minutes, though few would enjoy those ten minutes.

However, suppose that the question was to find all the integer roots of 10,604,499,373? (i.e., solve a^b= 10,604,499,373). Without knowing either 13 or 9, that would be a very long calculation to find an "a" and "b."You would have to use trial and error, starting with 3, 7, 9, 11, and 13 and multiplying them out until you get 10,604,499,373. It could be done with a pencil and pad—just not in 10 minutes unless you were some sort of math savant.

X= a^b is a problem that is hard to solve if you know one side of the equation but easy if you know the other. It is easy if you know a^b but hard if you know X. On the other hand, if you are a computer that can plug in numbers very rapidly, it may not be too hard to solve the problem only knowing X.

Fortunately for cryptography, there are several types of problems easy to solve in one direction but hard—*even for computers*—to solve the other way.[36] And public-private key encryption typically employs two types of these problems.

First, "modular exponentiation," your child in elementary school would call "clock math." In modular arithmetic, you add numbers within a limited group of integers, i.e., like a clock. In clock arithmetic, modulus 12 or "mod 12," 7 + 6=1 (not 13 because there is no 13 on a clock dial). Of course, you can choose any number for your modulus. If you choose very big prime numbers, things get computationally burdensome.

Exponential modulus math creates a unidirectional hard problem.

$$m^e mod\ N \equiv c$$

Easy →

This equation could be used for cryptographic purposes. *M* is a fixed number in some cipher, analogous to "3" in the above example of the Caesar code. It is raised to the *e* power (*e* being a public, random number) and then divided by *N*—the "modulus," again a public number. I.e., 4^4 mod 12 = 4. This operation produces *c*. It is easy to find *C* if you know *m*, but *not* vice versa:

$$?^e mod\ N \equiv c$$

← Hard

Thus, the challenge for cryptography. Find a "key"—the secret number that takes us from *c* back to *m* and avoid the impossibly hard calculation. The "key" can be expressed as follows:

$$c^{d} mod\ N \equiv m$$

The *d* is the key—the "decryption" value—that takes you from *c* to *m* without doing the impossibly burdensome calculation. And, it must be chosen in a way so that you cannot figure out *d* from *e*—though they are, in a sense, inverses. That means you need *another* unidirectionally difficult equation. And, that equation involves prime factorization and a function that relies on it: Euler's Φ function.

It was Euclid who discovered all numbers can be reduced into a product of prime numbers, a/k/a "the prime factorization." 12 reduces to 3*4; 17 is reducible to 17*1.The Swiss mathematician, Euler, took this idea and created the Φ function. It looks at all integers less than a given number and counts how many of these integers have *no* common prime number in their factorization. Thus, $\Phi[8] = 4$ because the number 8 shares no common prime factors with 1, 3, 5 or 7.[37] Prime factorization, necessary to figure out the Φ function is burdensome, even for computers.[38] However, Φ function is not terribly burdensome for *prime* numbers. Since by definition, a prime number is not divisible by any integer except itself, *no* integers less than a prime number have a common factor with the prime number. $\Phi[8] =$ [7, 6, 5, 4, 3, 2, 1] = [6]. Thus, for any prime number $\Phi[x] = x - 1$.

You can use Euler's Φ function to make a problem that is *really* hard in one direction. Take any two prime numbers, P1 and P2. Multiply them together and get *N*. What is $\Phi[N]$? As mentioned above, you could chug through all the prime factoring—which, if the number were large enough, would take centuries for a computer. However, *you* know the answer because you know the two prime numbers—and if you do, then $\Phi[N]$ is easy to figure out due to the algebraic features of Euler's Φ function.[39] Thus, P1 and P2 are, in a sense, your private key. Their product is the public key. It would take forever for a computer to figure out the prime factorization.
The question is how to apply Euler's Φ function to cryptography to obtain two numbers that are multiplicative inverses but computationally impossible to derive from each one-way, i.e., a public and private key.

Now, the trick is to combine the features of Euler's function with modular exponentiation.
It is Euler's theorem that connects his function to modular exponentiation. This following relationship holds for any two numbers without a common factor—you will always get one.

$$m^{\Phi(n)} \equiv 1\ mod\ n$$

This formula reduces with the application of Euler's theorem so as to define d as $\frac{k*\Phi(n)+1}{e}$.[40]

Returning to the original encryption using modular exponentiation:

$$m^{e} mod\ N \equiv c$$

$$c^d mod\ N \equiv m$$

Now, you can select your d and e in way that follows this equation: $d = \frac{\Phi(n)+1}{e}$. Observe that is very difficult to figure out d unless you know the prime factorization of n, *i.e.,* the two prime numbers that make up N and which make the Euler's function $\Phi(n)$ easy to solve. D is the private key, and e is the public key—and you can't derive e from D without knowing $\Phi(n)$—but if you know D, you can easily derive e.

Math Problems that are Hard in Only One Direction and Digital Signatures

This principle is used in digital signatures, which are used to authenticate identity. In digital signatures, the Signer creates a public and private key. He submits the public key to a certifying company like GlobalSign, DigiCert, or GoDaddy. The company examines identifying documents discussed above, then posts the public key on the web and issues a certificate to the Signer with information about the Signer, the authority, and the public key.

The Signer then encrypts a document and sends it to the Receiver with the certificate. The Signer also makes a "hash" of his document using a publicly available algorithm. A hash is a mathematical transformation of a text. Unlike public-private keys, a hash is a one-way function. It is a *unique* number that an algorithm identifies from a particular string of text. The Signer uses his private key to encrypt the hash. He sends the Receiver the certificate, the document, and the document hash.

The Receiver inspects the certificate and obtains the public key from the certifying company. The Receiver then hashes the document the Signer sent using the same public algorithm, obtaining a hash value. The Receiver, using the public key from the certifying company, then decrypts the encrypted document the Signer sent. If the two hashes have matching values, then the Receiver can be confident the Signer is who he purports to be.

ZDKs and Math Problems that are Hard in Only One Direction

Applications using zero-knowledge proofs also take advantage of these unidirectional computation-heavy problems. They allow a user (a/k/a Prover) to use a mathematical transformation of a hash of some "secret." The secret could be the Prover's age demonstrated to a trusted third party using one of the methods described above. In other words, the secret could be entrusted to a third entity.

The Prover would be able to reveal a value of mathematical transformation of that secret into a public value—from which it would be near impossible for the Verifier to figure out the secret. In other words, the Verifier's "secret" would be "x" in the following now familiar equation from modular exponentiation:

$$g^x \text{ mod } P = \text{public key}$$

Easy

$g^x \bmod P = \text{public key}$

Hard

The Verifier could not guess "x" from the "public key" if the numbers were large primes; it would be too computationally difficult.

The platform (a/k/a Verifier) would then present the Prover with a math problem that could only be solved *if* the Prover knew the value of "x". But, just like the Prover in the cave, the Prover would never reveal the value of "x" to the verifier. Nor would the Prover reveal anything else about himself.

Fiat-Shamir was one of the first and most influential mathematical protocols to operationalize these insights.[41] And, since then, numerous mathematical and computer techniques have developed.[42] The outlines of how this proof works as follows:

The Prover (the internet user that wants to create an account with the platform (Verifier)) demonstrates that he is over 18, generates a random integer "r" which he uses to generate C in the following way: $C= g^r \bmod p$. The Prover keeps r private and reveals C to the Verifier. (Again, due to the unidirectional difficulty of the modulus exponentiation, the Verifier cannot figure out r .)

The Prover also has a secret value x. It could be a hash value encoding his age, citizenship, or car insurance. It would be obtained from a trusted third party online—just like the certificate for a digital signature. The trusted third party would also provide the Verifier with the information needed to perform the challenge, i.e, the public values of g and p. The Prover calculates: $y= g^x \bmod p$—and reveals Y to the Verifier. (Again, due to the unidirectonal difficulty of the modulus exponentiation, the Verifier cannot figure out x .)

The Verifier now knows C, g, p and y. The Verifier then presents a challenge—analogous to asking the Prover to appear at the north side of the cave or asking whether he switched the colored balls. The challenge the Verifier asks: what is w, where $w = x + r \bmod (p-1)$?

Because the Prover knows x and r, he can easily respond with the value "w." (All the other values are public.)

While the Verifier does not x or r, he knows the following relationship (if he remembers how to do algebra with exponents and modulus arithmetic[43]:

$$y*C \bmod p=g^w \bmod p$$

Notice all of these values are public—but that the relationship will only hold *if* the Prover knows the correct w, which can only be calculated if the Prover knows the secret "x." Thus, if this relationship holds, the Verifier can be confident the Prover knows x, without it every being revealed.

There is, however, an easy way to cheat. The Prover could choose a random value z, and it could select C so that it equals $g^z/y \bmod p$. When the verifier asks for the value of W, the Prover sends in Z. This answer will satisfy the relationship: $y*C \bmod p = g^w \bmod p$.[44]

To avoid that problem, the Verifier can randomly ask for r or w (recall w = x + r). If the Verifier asks for x + r, then the fraudulent Prover can return an W based on $g^z/y \bmod p = C$. But, then the fraudulent Prover could not answer what r is. (Remember C is defined as $g^r \bmod p$. You cannot figure out r from C!). Similarly, if the Verifier asks for r, the fraudulent Prover could provide it by defining C as $g^r \bmod p$. But, if the fraudulent Prover did that, then he cannot define C as $g^z/y \bmod p$.

Blind Signatures. Another cryptographic technique that has application to age verification and authentication more broadly is blind signatures. First developed by David Chaum,[45] this technique allows banks to verify cash purchases or governments to verify voting, without the verifier (bank or government) knowing what the cash was spent on or for whom the voter voted.

To use an analogy, consider a ballot that a particular official needs to verify as being cast by an eligible voter. He casts his ballot and places it in a special carbon paper envelope. He then places the carbon paper envelope in a normal paper envelope, with his address written on outside of the envelope, and sends it to the ballot checker. She checks his name on her voter roll and makes sure he has not voted before. Then she opens the normal paper envelope and signs her name on the outside of the carbon envelope—and sends it back to the voter. The signature transfers through the carbon paper to the ballot. The voter's ballot now has the verifier's signature, but the verifier *never* saw the contents of the ballot.

This trick is achieved through various applications of public and private keys, elaborating on the feature that some math problems are easy to solve in one direction, allowing easy encryption but difficult decryption.

Application to Age Verification. ZKP could be used for age verification. A user seeking to establish his identity could establish his or her identity with a trusted third party using any of the methods outlined above. Once age has been verified using any of these methods, the third party would permanently delete all records submitted to authenticate age.

The trusted third party would provide the user/ Prover with a number (or a hash of a number) and would agree with the social media firm upon a particular type of problem (the "challenge"). The social media firm—now acting as Verifier-- would then present the user/Prover with the problem. Because you have the number obtained from the third party, you would then solve the problem using the number. The Prover could then show that he possesses the "number" received from the

Trusted Third Party. But, the Prover would never reveal the number or anything about your identity—*except that you are above the prescribed age*.

Application Beyond Age Verification: Proof of Humanity. With artificial intelligence, human authenticity is bound to emerge as a central problem for internet usage. Without confidence that internet users are dealing with human beings, humans will trust the internet to a lesser degree. However, this problem should not result in one large registry of human beings, which would no doubt be subject to abuse by whoever controlled it. Instead, ZKP and similar approaches can demonstrate human authenticity without revealing anything else about the authenticated human.

Conclusion. As AI makes it ever more difficult to know if you are dealing with a human online, authentication will become essential. Many are already proposing a universal registry of all human beings on the planet using biometric information such as retinal scans, a frightening thought from a privacy perspective. But, it is possible to fight for both privacy and age verification by adoption of techniques described here, such as blind digital signatures and zero-knowledge proofs. These age verification approaches protect children and privacy and eliminate burdens on First Amendment-protected adult speech.

Endnotes

1. Arkansas Sess. Stat, Act 614, To Create The Protection Of Minors From Distribution Of Harmful Material Act; To Establish Liability For The Publication Or Distribution Of Material Harmful To Minors On The Internet; And To Require Reasonable Age Verification (2023); Louisiana H.B. 142, 2022 Sess. 440, To enact R.S. 9:2800.28, relative to material harmful to minors; to provide for liability for 3 the publishing or distribution of material harmful to minors on the internet; to 4 provide for reasonable age verification; Texas H.B. No. 1181, An Act relating to restricting access to sexual material harmful to minors on an Internet website; providing a civil penalty; Virginia S.B. 1515, Harmful materials; civil liability for publishing or distributing to minors on the Internet.
2. See generally Clare Morell et al., Protecting Teems From Big Tech: Five Policy Ideas for States (EPPC & IFS, Aug. 2023).
3. Brett Burstein et al., Suicidal attempts and ideation among children and adolescents in US emergency departments, 2007-2015, 173.6 JAMA Pediatrics 598 (2019); Katherine M. Keyes et al., Recent increases in depressive symptoms among U.S. adolescents: Trends from 1991 to 2018, 54 Social Psychiatry & Psychiatric Epidemiology 987 (2019); Melissa C. Mercado et al., Trends in emergency department visits for nonfatal self-inflicted injuries among youth aged 10 to 24 years in the United States, 2001-2015 318.19 JAMA 1931 (2016); Ramin Mojtabai et al., National trends in the prevalence and treatment of depression in adolescents and young adults, 138.6 Pediatrics (2014); Gregory Plemmons et al., Hospitalization for suicide ideation or attempt: 2008-2015, 141.6 Pediatrics (2018); Jean M. Twenge et al., Age, period, and cohort trends in mood disorder indicators and suicide-related outcomes in a nationally representative dataset, 2005-201, 128(3) J. Abnormal Psychology 185 (2019).
4. Jean M. Twenge et al., Worldwide increases in adolescent loneliness, 93 J. of Adolescence 257 (2019).
5. Jean M. Twenge et al., Trends in U.S. adolescents' media use, 1976-2016: The rise of digital media, the decline of TV, and the (near) demise of print, 8(4) Psychology of Popular Media Culture 329 (2019).
6. Joan Luby, Increasing suicide rates in early adolescent girls in the United States and the equalization of sex disparity in suicide: The need to investigate the role of social media, 2(5) JAMA Open e193916-e193916; Henry A. Spiller, et al., Sex- and age-specific increases in suicide attempts by self-poisoning in the United States among youth and young adults from 2000 to 2018, 210 J. of Pediatrics (2019).
7. Hunt Allcott, The welfare effects of social media," 110(3) Am. Econ. Rev. 629 (2020).
8. Pierce v. Soc'y of the Sisters of the Holy Names of Jesus & Mary, 268 U.S. 510, 535 (1925).
9. Troxel v. Granville, 530 U.S. 57 (2000).
10. F.C.C. v. Pacifica Found., 438 U.S. 726 (1978).
11. 47 C.F.R. § 73.3999.
12. 397 U.S. 728 (1970).
13. Ashcroft v. American Civil Liberties Union, 542 U.S. 656 (2004).
14. *Id.* at 666-67 (2004).
15. § 10:1. The right of a child to make contracts and disaffirm them, 1 Leg. Rts. Child. Rev. 3D § 10:1 (3d ed.); R.T. Bottiger, Comment, Infants' Contracts and Their Enforcement, 35 Wash. L. Rev. & St. B.J. 465 (1960).
16. E. Allan Farnsworth, Contracts § 4.5 (4th ed. 2004).
17. A.V. v. i-Paradigms, Limited Liability Company., 544 F.Supp.2d. 473 (E. D. Vir. 2008), aff'd 2009 WL 1015145 (4th Cir. 2009); E.K.D. ex rel. Dawes v. Facebook, Inc., 885 F. Supp. 2d 894 (S.D. Ill. 2012)
18. See, e.g., S.D. Codified Laws § 26-10-19 ("No minor may be tattooed unles s the minor's parents have signed a consent form authorizing the tattoo.").
19. See, e.g., Parham v. J.R., 442 U.S. 584, 604 (1979) ("The fact that a child may balk at hospitalization or complain about a parental refusal to provide cosmetic surgery does not diminish the parents' authority to decide what is best for the child."); Kanuszewski v. Michigan Dep't of Health & Hum. Servs., 927 F.3d 396, 418 (6th Cir. 2019) ("Parents possess a fundamental right to make decisions concerning the medical care of their children.").
20. See, e.g., Mass. Gen. Laws Ann. Ch. 138, § 34; Wis. Stat. Ann. § 765.02.
21. See, e.g., Al. Veh. Code § 12650(b); Cal. Veh. Code § 12650(b); Fla. Stat. Ann. § 322.09(1)(a).
22. 10 U.S.C. § 505 (1983).
23. See, e.g., Mass. Gen. Laws Ann., ch. 119, § 55A (waiver of counsel by minor in juvenile delinquency proceedings must be made through parent or guardian).
24. Cal. Bus & Prof. Code §22706 (prohibiting persons under the age of eighteen from using artificial tanning devices without parental consent).
25. Turner Broad. Sys., Inc. v. F.C.C., 512 U.S. 622, 662 (1994).

26. Ashcroft v. Am. C.L. Union, 542 U.S. 656, 667 (2004).
27. John Ehrett & Clare Morell, Age Verification: Policy Ideas for States (IFS & EPPC, 2023).
28. Department for Science, Innovation and Technology and Paul Scully MP. Press release, Online Safety Bill bolstered to better protect children and empower adults (June 30, 2023), https://tinyurl.com/32fx64vz.
29. Online Safety Bill, Part 3— Providers of regulated user-to-user services and regulated search services: duties of care Chapter 2 — Providers of user-to-user services: duties of car, https://bills.parliament.uk/publications/51870/documents/3679
30. 542 U.S. 656 (2004).
31. Sara Ahadi & Andrew Carroll, Developing an aging clock using deep learning on retinal images, Google blog, (Apr. 2023), https://ai.googleblog.com/2023/04/developing-aging-clock-using-deep.html
32. Department for Science, Innovation and Technology and Paul Scully MP. Press release, Online Safety Bill bolstered to better protect children and empower adults (June 30, 2023), https://tinyurl.com/32fx64vz.
33. https://www.aleo.org/
34. https://medium.com/@aleohq/practical-use-cases-for-zero-knowledge-d8e5be5dfe46.
35. Jean-Jacques Quisquater, et al., How to Explain Zero-Knowledge Protocols to Your Children
36. I am indebted to the excellent YouTuber "Art of the Problem," https://www.youtube.com/watch?v=wXB-V_Keiu8.
37. In contrast, the number 8 has a common factor of 2 with 6, of 2 and 4 with 4, and of 2 with 2.
38. Factoring large enough prime numbers can take computers years, even centuries.
39. For the interested,
40. Here's the derivation for the curious:
41. Nalin Bhardwaj, A Succinct Story of Zero Knowledge, https://nibnalin.me/assets/zk.pdf
42. I am indebted to the superb YouTuber, "Alex on Science." His tutorial on zero knowledge proofs is wonderfully accessible. https://www.youtube.com/watch?v=cI5lkif-V1c.
43. For the curious, here's how to solve the problem:
44. For those desirous to know why, the algebra is simple: $Y^* C \bmod p = yg^z/y \bmod p = g^z \bmod p$.
45. David Chaum, Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms, 24(2) Communications of the ACM 84 (1981).