DEMOS

REWIRING THE WEB

THE FUTURE OF PERSONAL DATA

JON NASH
CHARLIE HARRY SMITH

JUNE 2023

Published by Demos June 2023
© Demos. Some rights reserved.
15 Whitehall, London, SW1A 2DD
T: 020 3878 3955
hello@demos.co.uk
www.demos.co.uk

CONTENTS

ACKNOWLEDGEMENTS

We would like to thank a number of people for their input and feedback on the draft of this paper. These include: Keegan McBride, Vicky Nash, Roger Taylor, John Taysom, Kirsty Innes, Tom Westgarth, Leo Ringer, Andrew Bennett, Simon Worthington, Mark Bembridge, Areeq Chowdhury, June Brawner, Elliot Jones, Valentina Pavel, Will Hayter, Bluebell Drummond, Vincenzo Rampulla, Casey Calista, Samuel Rowe, Bryan Glick, Dave Birch, and Ian Brown.

At Demos, we would like to acknowledge the assistance of Polly Curtis, Ellen Judson, and Oliver Marsh throughout the project.

ABOUT THIS PROVOCATION PAPER

This paper proposes a series of technical, regulatory, and institutional interventions that reimagine the foundations of a modern internet built on privacy, interoperability, and consent.

Jon Nash
Charlie Harry Smith

June 2023

JON NASH

Jon is a political scientist and entrepreneur. He co-founded Mainstream, a live video streaming platform, with former Facebook and ABC executives, and went on to build a popular location based messaging service. His research focuses on technology policy and democratic innovation and he has advised the goverment on replacing General Data Protection Regulation and the future role of the ICO.

CHARLIE HARRY SMITH

Charlie is a political philosopher and doctoral student at the Oxford Internet Institute. His research considers the normative and theoretical issues surrounding digital identity systems, with a particular focus on contemporary governmental policy in England and Wales. Charlie also regularly consults on global identity projects, and currently advises the Open Identity Exchange, the trade body for UK digital identity companies.

EXECUTIVE SUMMARY

In this paper, we argue that the widespread use of personal information online represents a fundamental flaw in our digital infrastructure that enables staggeringly high levels of fraud, undermines our right to privacy, and limits competition.

We present an alternative system where standardised requests are instead routed by a user's device, with their consent, between certified organisations. This allows their personal information to be substituted for secure alternatives, like unique identifiers, claims, and tokens.

For example, an online retailer could make a request for 'payment' instead of asking a customer for their card details. The user's device would then match this request to the organisations that could respond and present these options to them in a standardised consent dialogue. Once selected, the payment request would be forwarded by the user's device to their bank, which would respond directly to the retailer with a one time payment token that only they could use.

The ability to securely move information between trusted organisations—with user consent—would have a profound effect on all aspects of the web. In particular, we explore how digital identity, online payments, and digital advertising would be affected, and describe the benefits of this system for both users and organisations.

Finally we argue that the common carrier laws that already apply to internet service providers should be extended to our devices and the routing of standardised requests. That a new national certification authority is needed to establish trust and resolve liability, and that standards for requests and responses should be set in cooperation with existing standards bodies and consortia.

Together, these technical, regulatory, and institutional interventions reimagine the foundations of a modern internet built on privacy, interoperability, and consent.

INTRODUCTION

The web's creators did not set out to build the foundations of our twenty-first century economies. They could never have predicted the volume and variety of services the web would one day handle. What started life as a communications tool for academic and military researchers now lets us do almost anything, from shopping for groceries to applying for a mortgage. But performing these tasks today involves the use of large amounts of personal information. We are constantly expected to acquire, remember, and provide information about and relating to ourselves when interacting with organisations; not just usernames and passwords, but bank account numbers, addresses, national insurance numbers, and even doctors' letters and utility bills.

The web has catalysed huge levels of growth and innovation, but our approach to personal information has become not just a bottleneck, but a liability. Managing all this information now limits everything from our access to government services to the health of our democracy. With social networks struggling to distinguish humans from bots, bad actors can influence the public discourse on a massive scale. At the same time, safely making payments, providing our details, and proving who we are is becoming ever-more challenging. And, against this background, the usability of the web has steadily declined.

Indeed, our continued reliance on personal information is fuelling a security and privacy nightmare: as many as 82% of all data breaches today stem from the misuse of credentials[1]. Behind the scenes, companies and governments are struggling to keep up. In the perpetual arms race to protect our personal information, the criminals are winning—as the Reverend Mike Hall discovered in 2021. Hall returned home after a few weeks away to find his belongings gone, someone else living in his house, and new building work underway[2]. It turned out a fraudster had used a fake driver's licence to set up a bank account in Hall's name before selling his home from underneath him. The new owner, who had legally bought the property from the man he thought was Hall, was none the wiser.

Although extreme, Hall's story illustrates both how brazenly fraudsters are profiting from the status quo, as well as just how dramatically the use of personal information—and particularly our credentials—is failing us. To take another example, Britain was last year crowned the card fraud capital of Europe, with 84% of attacks using stolen card details[3]. Yet because we reuse the same payment details everywhere we shop, if these credentials ever do get into the wrong hands we have to throw them away and start again, waiting for our sensitive banking details to be posted to us on another plastic card.

To realise a web fit for the twenty-first century, we need to fundamentally rethink the ways in which we interact with organisations online. We must look beyond the personal information that fuels fraud and adds friction, and challenge the idea that we should be personally responsible for remembering, managing, and repeatedly entering all this information ourselves.

In this paper, we propose a set of technical, regulatory, and institutional interventions that would realise a web built not on personal information, but on ***trusted connections***. An important insight underpins this proposal: if the right organisations could ask the right questions of one another, then our information could get from where it is to where it needs to be without us having to read it out, write it down, or type it in. This ability—to reliably ask for and provide data—is therefore key to making the web faster, safer, and more usable.

[1] https://www.verizon.com/business/resources/reports/dbir/2022/master-guide/
[2] https://www.bbc.co.uk/news/uk-england-essex-59069662
[3] https://www.smf.co.uk/uk-is-card-fraud-capital-of-europe-think-tank/

REPLACING PERSONAL DATA WITH TRUSTED CONNECTONS

Today, if a company wants to contact us, they ask for our email address. To take payment, they ask for our card details. And, to sign us in, they ask for our username and password. The system we propose is radically different. It would allow us to do these things—and many more—by creating trusted connections between existing organisations, without having to share any personal information.

Take online payments as an example. Instead of typing your long card number, expiry date, security code, full name, and home address into a retailers website, a request for payment would be routed, by your device, to your bank. Your bank would then be able to respond directly to the retailer with a unique payment token that allowed the payment to be made. While this describes one example, the same model would apply to almost every interaction we have online.

This would all be enabled by your device, which would build up a list of who had what, functioning as a private directory of the organisations that you interact with. When another organisation needed to know something, it would simply ask for it in the form of a specific request. Your device would then route these requests to the relevant organisations, who would each respond directly with the appropriate information[4].

Importantly, however, no connections would ever be made without your devices first securing your ***explicit consent***. Meaningful consent is currently hard to come by on the web. We are regularly faced with so many policies, terms, and conditions that we can do little more than blindly agree. Compounding the problem, user interfaces are often designed to maximise click-through rates to service the interests of organisations rather than users.

Instead of giving each company latitude to ask for 'consent' in their own way, on sites and in-apps, the same standardised screen would be used across

FIGURE 1
DIAGRAM OF THE SYSTEM ARCHITECTURE

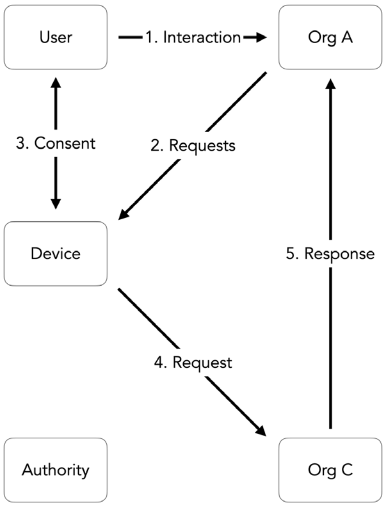

[4] This describes what we call a dynamic request as it creates a direct connection between two organisations, but in some cases, blind requests could be made that would route the response back through the device. This would allow us to share information without revealing the origin of the request.

FIGURE 2
MOCKUP OF THE CONSENT INTERFACE

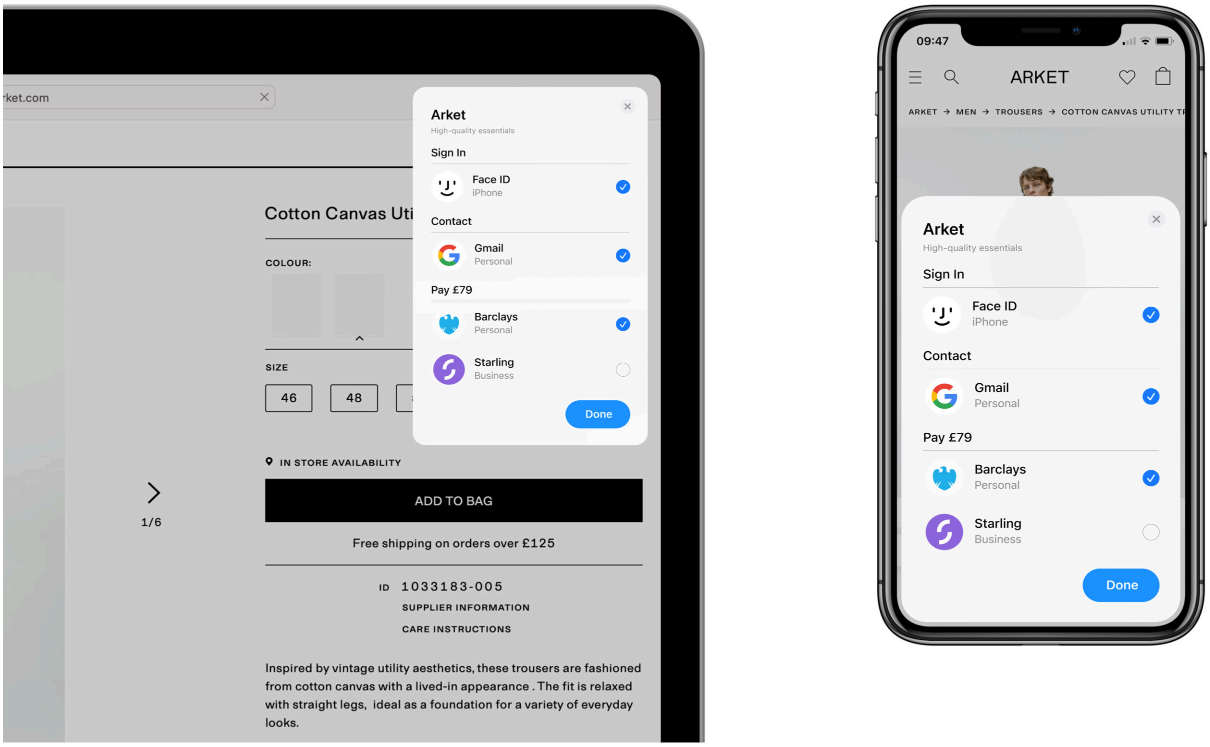

Inspired by vintage utility aesthetics, these trousers are fashioned from cotton canvas with a lived-in appearance . The fit is relaxed with straight legs, ideal as a foundation for a variety of everyday looks.

different devices and manufacturers. From a user's perspective, giving consent would be transformed into a consistent process. Their device would clearly show three things: the organisation making the requests, the types of requests being made, and the names of the organisations or services in the user's life that could respond. They would then be able to make an informed decision and better understand who had their information.

This would always be extremely straightforward for users. They would not have to set anything up, and their device would never redirect them to a browser or authenticator, ask them to enter any personal information, or accept extensive terms and conditions.

Of course, we would want to know that the companies asking for our information had some legitimacy; that they were not trying to defraud us, steal, or sell our information. Likewise, organisations requesting information would need to know that it was coming from a legitimate source; that it could be relied upon, and would not expose them to undue risk. For this reason, we argue that a recognised authority should be established to set and certify the requirements for different types of requests within this system.

For low-risk interactions, minimal requirements would be set, but we would expect tougher requirements to be put in place for organisations requesting or providing more sensitive information. Taking out

FIGURE 3
DIAGRAM OF THE DEVICE QUERYING THE RECORD OF CERTIFIED ORGANISATIONS

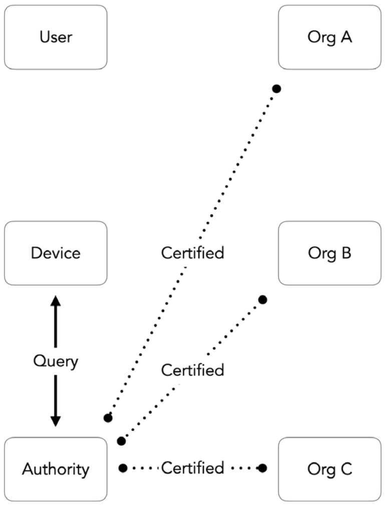

a mortgage, for instance, would require higher levels of assurance than subscribing to a streaming service. There is precedent here. In many sectors, like banking and aviation, we already expect governments to guarantee a level of protection by licensing or certifying companies to act. It is not, after all, left up to consumers to audit the liquidity of banks or assess the safety of airlines and we think the same model should be applied to our information.

The benefits of ensuring that participating organisations were certified to handle our information would be manifold. Going forwards, individuals would know that any interactions handled in this way would always be coming from or going to trusted organisations. This would massively reduce the risk of phishing attacks, scams, and financial fraud, removing the burden on users to check a site's SSL certificate or URL, and thereby making it far more difficult for them to mistakenly give their details or data to malicious actors.

In much the same way, organisations would also be able to interact with greater confidence, knowing they could trust those that they were interacting with. But certification would save participating organisations a considerable amount of time and money, too, as they would take on significantly less liability when sharing or accessing data from certified entities. Additionally, to begin operating within this system, all these organisations would need to do is become certified. This low barrier to entry, coupled with the reduction in liability, would therefore be extremely appealing.

Certification would also realise a powerful governance mechanism, helping ensure sufficient oversight and accountability. Each country's certification authority would, for instance, be able to revoke certification if organisations misbehaved. We would also expect regular auditing to accompany higher levels of assurance. While similar processes to these already exist on parts of the web, we think such decisions should be handled by public bodies embedded in the legal and political framework of each country—not the private companies that currently provide unaccountable accreditation and certification functions.

The last major aspect of this proposal involves ***standardising*** requests. To facilitate easy and secure connections between trusted, certified organisations, everyone would need to speak the same language. At the moment, any organisation looking to interoperate with others on the web first has to register and integrate with separate data providers, as each maintains their own bespoke application programming interfaces (APIs).[5] When there are only a few information providers in an ecosystem, this is not necessarily a problem. Organisations are generally happy to spend development time integrating with each organisation and accept each providers' governance demands. But this proprietary approach quickly becomes unworkable at scale. Organisations providing access to information become overburdened, while smaller organisations are left structurally disadvantaged. The result is hugely damaging for competition and innovation.

By contrast, in this model, lots of different requests could be routed between lots of different organisations. The way in which each organisation asked for or provided information would therefore need to be standardised. This would ensure that all actors in the system could seamlessly interoperate with one another, and is key to realising the benefits of an open, flexible ecosystem built on a foundation of trusted connections and certified organisations.

Of course, in many sectors, like telecommunications and banking, industry participants already develop and maintain standards via various international organisations and consortia. Standards setting would accordingly be largely left to these organisations. But the outputs of these bodies would need to be consolidated into a unified record, published by a new international organisation—a ***standards forum***, rather than a standard-setting body. We believe that this forum should also include a 'layer' of civil society organisations, to advocate for the rights of citizens and counterbalance industry interests in the standardisation process.

Setting universal standards would usher in numerous advantages. It would save time, reduce costs, and enable a much higher volume of interactions to flow through the system, ensuring that organisations knew what information to expect as well as how to handle requests and responses. Integrating with any potential organisation would become far more straightforward, opening up possibilities for innovative new use cases.

The transformation would be analogous to that which revolutionised the railways. Before standardisation, different rail companies used different gauges of track. Mandating a standard gauge enabled these tracks to interconnect and kick-started the technology's massive expansion. For similar reasons, there would be a clear incentive for organisations to use the agreed upon standards. Following these standards would be a prerequisite for certification, which in turn would grant organisations access to this system, the benefits of which we discuss in the following chapter.

[5] APIs function somewhat like pipes, connecting software at two organisations together through a datastream.

PRESERVING PRIVACY

Rebuilding the web on a foundation of trusted connections would realise numerous advantages. Not only would the important interactions in our lives become more secure, easier to make, and based on our explicit consent, but this shift would also open up new opportunities for interactions that are not possible today.

We have already seen how managing personal information exposes us to significant privacy and security risks. If and when our information is compromised, a single breach quickly becomes a catastrophe, as the effects of that breach cascade through all the different contexts in which we have previously and repeatedly entered our data. The reuse of personal information therefore magnifies the chances of, and negative impacts of, its misuse. This is bad enough. But as well as introducing such structural weaknesses to the web's foundations, expecting individuals to manage their own personal information also allows them to be tracked and profiled across these various contexts.

In fact, a whole industry of advertisers and data brokers, some more legitimate than others, currently profits from the processing of personal data. Tracking our digital footprints, these companies build up detailed profiles of our interests, which they then resell or else monetise—undermining our privacy. What is worse, successive attempts to bring these companies to heel via data protection regulation has done little to curb their appetite for information, indicating again just how broken the system has become.

We think the root cause of all these issues is a lack of specificity. Namely, the personal information that we currently replicate and reuse all across the web is neither context nor function specific. Your email address, for example, can be used by anyone, to send you anything, at any time—and it is associated with all of your accounts. Once any of this information gets leaked, stolen, or sold, it can therefore quickly be put to work (against you) in another context, to realise functions that you did not originally intend and to which you did not consent.

In this way, data about you is somewhat analogous to nuclear waste[6]—valuable if it can be processed in well-managed, high-security facilities, but dangerous if improperly handled or, worse, allowed to leak out into the environment. Our proposal takes this provocation seriously, addressing the twin issues of context and function imprecision by building in hard limits to the connections we make. This amounts to a radical reimagination of the role of personal information on the web. By ensuring that users must explicitly consent to any connection, and locking in these hard limits as part of the standardisation process, we reduce the radioactivity, ensuring that, if data does leak, it cannot cause the widespread damage that personal information invites.

To understand how this proposal moves us towards a more private and secure system, based on notions of context and function specificity, there are three different technical elements that each protect a user's privacy and keep their data safe. These are unique identifiers, tokens, and claims. Together, these three elements would allow us to achieve much of the functionality of the current web, simultaneously unlocking entirely new possibilities, while eliminating the morass of unspecified personal information that currently limits our online interactions.

UNIQUE IDENTIFIERS

The use of unique identifiers would dramatically change the way that organisations assessed who they were dealing with in online interactions. Currently, organisations store an email address and password when you first sign up, then ask you to provide this information again when you next interact with them. This indicates that you are likely to be

[6] https://www.forbes.com/sites/johnkoetsier/2022/08/06/data-is-the-nuclear-waste-of-the-information-age-on-big-tech-and-privacy/

the same person. But, as we know, emails are easily copied or stolen, and many users do not choose secure passwords. Our use of personal information in this way therefore allows bad actors to commit fraud by posing as someone they are not.

The problem is that these identifiers are universal—they are the same across the many contexts in which we use them. This proposal would replace these universal identifiers with context-specific, pseudonymous identifiers[7]. Each of these identifiers would be one-of-a-kind, and only ever held by two parties. So, every organisation in your life would use a different random alphanumeric string to identify you, either via your device or else when communicating with other organisations directly.

These identifiers would still allow trusted connections to be made, but every time a device brokered a new relationship—either between itself and an organisation, or else directly between two organisations—a new, random identifier would be generated. When you first made a connection, the organisation involved would store this unique identifier instead of your email address and password. When you interacted again, your device would then automatically provide this unique identifier to reliably identify you. Indeed, users would not be able to see these identifiers which, following cybersecurity best practice, would consist of strings of randomly-generated letters and numbers that were always encrypted.

In the event of a data breach, the scope for negative repercussions would therefore be severely limited. Each identifier would not be a rich form of personal information; it would contain no sensitive details about you. And, behind the scenes, you would always be associated with a different identifier in each organisation's database. This means that the various entities in your life could not be linked up by bad actors, even if they did manage to acquire the unique identifier associated with your account in one particular context.

Unique identifiers, however, are only the first piece of the puzzle. In practice, they would rarely be exchanged on their own, and would mostly be accompanied by another element that contained the request or response necessary for an interaction to take place. These could take the form of either tokens or claims, which both build on the context-specificity of unique identifiers to designate a specific function or transfer a certain piece of information.

TOKENS

A token would allow an organisation to request some particular action, or respond to such a request. These requests and responses could be extremely function- and context-specific. To make a sign-in request, for instance, a one-time token would be sent alongside your unique identifier to instruct the organisation in question to log you in. Or, a contact token could specify that only three messages may be sent to the email associated with a particular unique identifier before that token expired. This specificity would be a powerful tool for ensuring that wide-ranging privileges were never granted to organisations, at least not without a user's explicit consent.

At the same time, tokens would also help guarantee a high level of security. For instance, if an organisation received a request without the appropriate token, or if the token did not correctly reference the appropriate identifier, then that organisation would ignore it. Because tokens would also be encrypted, only organisations with the relevant key could read the instructions they contained. The contrast with the status quo, where personal information is duplicated all over the internet and we can do little more than blindly trust that it will not be misused, would be stark.

CLAIMS

Rather than allowing something to happen, a claim would say something about us. They would usually be sent as a response to a request, and could take the form of a measure, such as a percentage or number, or simply a yes/no answer. Claims can therefore be far more privacy preserving than their personal information equivalents. For example, instead of providing your driver's licence to a car rental company, a request for licence confirmation could be routed to the driving authority. The authority could then send back a narrow response, specifying that you could drive, were over the age of 25, and had less than three points. In many cases, a simple 'yes' or 'no' response would suffice.

As this shows, the benefit of claims is that they allow organisations to say something about you without revealing significant amounts of personal information. Like tokens, they are functionally specific and constrained to one context—they respond to a single request and no more and, as they are also signed and encrypted, possess no value if intercepted.

BLIND VS. DYNAMIC REQUESTS

In some instances, users would want to conceal the

[7] In the industry, these are called pairwise identifiers. The IDs are kept safe on the device, protected by advanced biometric security features and encrypted on both your device and the company's servers.

origin of a request. You may not want your healthcare provider, for instance, to know that you were sharing health data with a particular organisation. In this case a 'blind' request would route the response back through the device instead of connecting the organisations directly. Blind requests would therefore facilitate a high level of privacy as the device would translate the unique identifier used with one organisation to the unique identifier used with another. The organisations involved would consequently not be told which other party made the request or responded.

While the use of unique identifiers, tokens, and claims would significantly improve users' privacy and security across all interactions, the greatest impact would likely be felt in situations that required identity assurance.

IDENTITY ASSURANCE CASE STUDY

Today, reliably assuring an individual's identity online is both challenging and costly. The problem is particularly acute in high-stakes environments, where fraudsters are highly motivated to circumvent security and defraud individuals and businesses. Given the well-known shortcomings of emails and passwords, developers have therefore been forced to add additional security mechanisms over time, continually ramping up security to try to stay one step ahead of criminals.

This is generally achieved by adding more authentication factors. These might involve checking something you know, such as a memorable answer or PIN, something you have, like a specific device or hardware key, or something you are, via a face or fingerprint scan. Combining several of these checks in a multi-factor authentication (MFA) process can help ensure that an organisation can trust the credentials they are being presented with actually belong to the person presenting them. As fraud gets increasingly sophisticated, however, these procedures have gradually become more byzantine, stacking up on top of one another to make sign-in a decidedly laborious task today in all but the lowest risk situations.

Arguably, companies have struggled to balance ensuring effective security with a good user experience. We all know how frustrating this can be. Most online interactions today begin with entering your email and password. You may then be asked for a memorable phrase but, more likely, the flow will be interrupted while you enter a code sent to you via SMS. In the most sensitive contexts, like accessing online banking and government services, organisations will then likely perform further probabilistic calculations in the background to assure your identity at even higher levels. This might involve checking whether your location is consistent with previous login attempts, or even recording how you are using your device to profile your unique behavioural signature.

As this shows, securely proving our identities has so far involved allowing organisations to analyse ever-more personal information that we provide in return for access to services. It has also required users to repeatedly wrap their heads around more and more steps, as each additional authentication factor is added—a trend that is likely to continue. But collecting all this data opens organisations up to significant risks while adding cost and complexity. This is not remotely sustainable. We need to take a different approach to proving identity on the web.

In many situations, it would be more efficient for our devices to handle basic identity assurance, leveraging the power of unique identifiers to build trusted and secure connections between the organisations already in our lives. Modern smartphones and computers, after all, already use advanced biometric technologies to ensure that only authorised users are granted access to a device. For most connections in a person's life, this process would provide more than enough confidence for a sign-in request to be made and responded to—and would completely replace emails and passwords with a far more secure and private alternative.

Some organisations, though, would need to check identities at higher levels of assurance. In these instances, your device would be able to respond to a request for a higher level of identity assurance by deriving a 'measure' of identity from the number, nature, and frequency of connections that had built up during normal use of your device. This information—akin to the sum of your digital footprint within the ecosystem—would provide an unprecedented measure of identity. The figure calculated by your device would take into account the entire map of your online presence without requiring any user input. But the added advantage of using this sort of measure instead of the personal information that companies currently collect and process themselves is that it would reveal almost nothing about you while still providing powerful assurances that you existed as an individual human being.

It would, after all, be extremely hard to fake this measure. Fraudsters would have to simulate a convincing, legitimate life within the ecosystem by building dummy connections with a large number of trusted organisations. This would be prohibitively costly and unrealistic, requiring them to generate months or even years of legitimate payments and interactions with certified companies under a false identity. Committing identity fraud would therefore become excessively difficult and expensive.

An even higher level of assurance may, however, be required when making the largest or riskiest transactions, such as a house purchase. In these cases, rather than asking the device for a measure of your identity, an organisation could go one step further and query the network of organisations in your life—banks, mobile providers, and government services—to generate a further measure of your digital footprint in aggregate across all their systems, too[8]. Faking this measure would not just be difficult, but next-to impossible. After all, fraudulently maintaining an entirely parallel life inside and outside of the ecosystem would require truly prohibitive levels of technical know-how and resources. And, once a fraudulent transaction was carried out, that device and the associated network of connections would be quickly 'burnt', requiring the fraudster to begin manufacturing another illegitimate identity from scratch.

If we think about it from an organisation's perspective: they would gain a robust way to get authoritative and secure answers to important questions—'is this person over eighteen?' and 'have I interacted with this user before?'—without, in many cases, needing to spend any money on developing or buying-in an identity assurance capability. Instead, all they would need to do is become certified to make sign-in requests under one, unified system that had the added benefit of supporting a wide range of other useful functions. Given that cyber crime and fraud cost the UK economy more than £4 billion last year, the cost savings from this alone would be astounding[9].

Simultaneously, organisations would gain the flexibility to ask for only what they needed, knowing they could trust these answers due to a rigorous certification and routing process. In many cases, this would mean they could opt to ask narrower, specific questions instead of collecting and storing significant amounts of sensitive personal information. While organisations requiring higher levels of identity assurance would still have the flexibility to collect additional data as required, this would save most actors in the system from being bound by the inefficiencies and risks associated with managing personal data. This would realise further cost savings while boosting user privacy.

What's more, it is worth pointing out that organisations would gain these benefits without needing to worry about falling foul of the legislation that governs user consent. GDPR, for instance, has strict guidelines around what counts as freely-given consent surrounding data capture and processing. But, under this proposal, the consent process would be standardised at the device level, not left to organisations. Unbundling the consent stage like this would mean that organisations would know that consent and certification were being handled by trustworthy, industry-wide processes that removed liabilities from their balance sheets.

From an individual's perspective, many of the same benefits would also be felt. First, they would save considerable time and effort thanks to a better user experience, all built upon a standardised process that they would recognise and trust across the web. In most cases, just one-click would suffice for proving who they were. Second, due to the certification process, they would know they were making connections with organisations they could trust. This would all make individual users more likely to create connections in the first place, without requiring users to manage their digital or physical identities themselves or rely on the services of third-party identity providers.

Third, and finally, individuals would make huge gains in terms of security and privacy, as they would no longer need to put their personal information at risk by reusing it across different contexts. This would stop shady organisations from building up a picture of how they acted across the web, limiting the cross-site tracking and profiling of users' actions—except where their consent had been explicitly and freely given. Not only would this achieve a level of identity assurance that is unimaginable today, but it would do so in a way that was extremely privacy preserving—and that, in many cases, would only require minimal user input.

[8] Both of these measures would use sector-wide, standardised formulas, debated and agreed upon in the standards bodies. This would ensure both democratic oversight and fair, secure levels of assurance.

[9] https://www.money.co.uk/credit-cards/fraud-report

ENSURING INTEROPERABILITY

Moving from a web built on personal data to a system of trusted connections would bring about a more open web, increasing interoperability and the scope for innovation. In other words, a standardised, regulated way to make and respond to requests would boost the capacity for flexible information exchange throughout the entire network. We would expect this to go some way towards undoing the entrenchment of various technology companies as the de facto gatekeepers of our online interactions.

Several safeguards would however need to be in place to fully realise the potential of genuine interoperability. In particular, routing requests, certifying organisations, and setting standards would all need to be done in an open, impartial way. No one company or organisation should be able to exert control over any of these key building blocks—not because we think technology companies are inherently anti-competitive, but because this is the best way to maximise competition and innovation in the system.

TRUSTED CONNECTIONS

As things stand, it is extremely difficult for all but the largest technology companies to make the data that they hold available to other organisations. They first need to build systems for registering and assessing developers, then must persuade individual organisations to integrate with their proprietary APIs. This presents a significant barrier to adoption, and makes it far more cost-effective for organisations to simply work with the biggest players.

Similar problems face organisations trying to access the data held by these gatekeepers. They must first register, often paying high fees, before spending a considerable amount of time and money integrating with each bespoke API. In other words, they must learn to speak a different language for each provider. This is not only exclusionary for smaller players; it is costly for the larger players, as they must maintain their own standards and associated governance processes. This is why we only ever see options to sign in or make payments offered by a handful of large technology companies that can absorb the costs necessary to develop such capabilities.

Under this proposal, organisations would no longer need to negotiate direct relationships with one another in this way. Interoperability would instead be achieved by leveraging our devices, which would build up and maintain a list of possible connections. Instead of requiring us to enter our information, any organisation could securely ask for or provide information in a standardised way. This would massively reduce friction for users and organisations, saving time and eliminating much of the possibility for error that currently exists.

This flexible system would allow market entrants to better compete with established players. After all, if a user had chosen to have a relationship with them, even the smallest organisations would show up on an equal footing. Additionally the barrier to entry for users to utilise any organisation's services would be extremely low. This would go a long way towards addressing the power and informational asymmetries that currently exist within the technology landscape—but could only be achieved by creating a highly interoperable ecosystem built on trusted connections.

CERTIFICATION AUTHORITY

Additional improvements in interoperability would arise from a standardised and open certification process. Rather than each responding organisation separately performing their own arbitrary assessment of the recipients of the data they held, certification would instead be handled by a new government body. This would reduce the level of liability that organisations currently shoulder, while also ensuring that users would have greater confidence in the connections being made within the system.

This would contrast dramatically with the status quo. Currently, decisions about access to data are predominantly made by unaccountable companies, on their own terms. Users accordingly have little control over how their data is used, not least as claims of 'legitimate interest' underpin extensive data processing. Comparatively, this proposal would open up a new mechanism for regulators to ensure that connections aligned closely with the interests and expectations of citizens. Both these aspects of certification can therefore be expected to increase interoperation across the network, and users' willingness to make connections in the first place.

STANDARDS FORUM

Finally, the proposed standards forum would allow sectors to define a ***lingua franca*** for making requests and receiving responses, maximising the scope for interoperability. This forum would work with existing standards bodies and trade organisations to develop and maintain these standards and publish a definitive record. Of course, in some cases standards may already exist. But, in most cases, we expect that a newfound ability to route secure, privacy preserving requests between organisations would lead to the development of new forms of standardised requests.

As we have already mentioned, the duplication of work required to integrate with various providers and their bespoke APIs currently proves hugely inefficient, disincentivising the development and adoption of innovative new solutions. Establishing agreed upon standards for requesting and providing data would therefore unburden all actors in the ecosystem, particularly allowing smaller organisations to more easily participate.

The result would look very different to what happens today, with larger players building their own APIs which encode their own values—often to the detriment of users and other organisations. Apple, for instance, has previously mandated that companies preference their sign in option above those of its competitors[10]. A unified language of standardised requests and responses would accordingly limit anti-competitive behaviour, help ensure an open, equitable ecosystem, and lay the groundwork for the creation of far more innovative products and services. To understand the value of interoperability, we can look at the online payments, and the role of online payment providers.

ONLINE PAYMENTS CASE STUDY

The Open Banking initiative demonstrates the UK finance industry is desperate to build safer, more direct relationships with businesses—especially following the failure of Paym, an ill-fated attempt to build a payments system around phone numbers rather than bank account details. As these schemes show, the industry knows it must innovate, and is looking to embrace a digital approach to improve its products.

Reusing banking details across the web evidently creates unneeded friction for individuals, who must manually re-enter their information whenever and wherever it is needed. But it also leads to staggeringly high levels of fraud, undermining user confidence. In 2021, fraud totalling £730.4 million was committed in the UK via cards, remote banking, and cheques, with Authorised Push Payment scams—where users are tricked into manually authorising transfers to criminals' bank accounts—accounting for an additional £583.2 million[11]. The need for a better mechanism for making payments could not be starker.

Sensing a market opportunity, a mass of companies are therefore trying to move the industry forwards. These intermediaries, from neo-banks to payments giants like PayPal and Venmo, promise a better user experience and more features than traditional financial services. In return, they generally extract fees from the transactions they process and generate further profits by monetising data they collect about customers. But the deeper problem, of course, is that none of these solutions actually tackles the root cause of the issue: our continued reliance on the manual entry of personal information.

Unsurprisingly, we do not think any of these approaches are desirable. A secure user-friendly process is needed that avoids the inadequacies of the traditional payments system and all its associated intermediaries. The system we propose would allow a request for payment to be routed by our device, with our consent, directly to our bank from a retailer. This would replace all the complication and fragmentation of modern checkout flows with a one-click process, standardised at the device level.

After pressing 'buy now', the retailer would simply send over a payment request. Querying the list of organisations, your device would then display your bank, and associated bank accounts—with your most recently used option likely pre-selected. Tapping to confirm, your device would handle authentication and forward the payment request. Your bank would then respond directly with a unique payment token, entirely removing the need for card numbers, expiry dates, and security codes.

Unlike universally-redeemable personal information, these tokens could be highly specific. They

[10] https://www.reuters.com/article/us-apple-apps/apple-asks-developers-to-place-its-login-button-above-google-facebook-idUSKCN1T6056
[11] https://www.ukfinance.org.uk/system/files/2022-06/Annual%20Fraud%20Report%202022_FINAL_.pdf

could, for instance, outline a specific one-time payment amount, or even define a series of recurring payments. This would be accompanied by information about the token's expiry and timing criteria. Optionally, a contact request token could also be included, allowing the retailer to send a receipt and, say, up to three marketing emails.

The difference for users would be immediately felt. Rather than entering personal information or trusting autofill tools with sensitive details, their device would instead make direct and trusted connections between the organisations they wished to interact with—all contingent upon their explicit consent, provided via a standardised, familiar, device-level interface. Except where they could offer genuine innovation worthy of user consent, the vast majority of payments intermediaries would therefore become obsolete under this proposal.

Individuals would accordingly save time on every checkout, gaining an unprecedented level of privacy and security along with a smooth user experience—no more frustrating one-time-passcode requests or in-app confirmations, and no more weak passwords reused across different sites. What is more, no personal information would be exchanged as part of this process, just unique identifiers and payment tokens. The knock on effect here would be that, as each connection would use a different identifier, it would be extremely hard for any third-party to track them across different interactions even if they wanted to.

But the real revolution would be for organisations. Fee-free, fraud-proof transactions would save companies and governments from swallowing the huge costs that currently accompany online payments. Take an online retailer turning over £10 million a year. Depending on merchant service fees, interchange fees, and their payment gateway, they could be paying as much as 4% on every transaction. These fees are levied by issuer banks, payments networks, and merchant banks, as well as actors like Apple, Google, PayPal, Worldpay and Square. Such staggering levels of intermediation could accordingly be costing our retailer as much as £400,000 a year in lost profits.

Under our proposals, that figure would drop to zero. There would be ***no fees*** for accepting payments made via a trusted connection between a merchant and the user's bank. Even for retailers with a physical presence, who may still have to cover the cost of physical point-of-sale equipment[12], the savings from being able to accept payments without intermediation would be astounding. And we would expect competitive companies to pass those savings along to their customers.

Of course, today intermediaries partly take a cut in order to build up a cushion to cover the cost of fraudulent transactions. Levying a fee allows them to invest in anti-fraud measures, but also foots the bill of facilitating refunds and chargebacks. This reflects the liabilities that payment providers and intermediaries shoulder—risks that are constantly fluctuating due to the ongoing arms race between fraudsters and legitimate businesses. By replacing card numbers, expiry dates, and security codes with payment tokens and unique identifiers, we would remove this rotten foundation from the financial sector, and could therefore expect a dramatic reduction in overall fraud levels.

After all, banks would be shouldering far fewer risks in the first place. This would render many of the costs associated with tackling fraud today moot. Transactions would only be made via direct, trusted connections between certified organisations. This would be highly secure[13] and would allow both parties to strengthen their relationships with customers. For all these reasons, we expect the financial sector to be a key vector for the adoption of these proposals.

[12] Our preferred solution would be to cover this functionality on the retailer's own device, but there could still be an incentive for retailers to use a dedicated POS terminal if it added valuable additional functionality.

[13] The scope for innovation here is huge. Users could create connections between their banks and third-parties, that would enable new use cases based on them leveraging purchasing data. Standards to support such connections would just need to be created in the appropriate bodies, then certified by the certification authority.

REQUIRING MEANINGFUL CONSENT

Along with greater privacy and interoperability, this proposal would also place greater emphasis on meaningful user consent. This is sorely needed. Under the status quo, users are often encouraged to choose actions they might neither have intended nor fully understood. ***Dark patterns*** exploit, amongst other tactics, the lack of a standardised design language for creating online interactions. And attempts by various regulators to reinforce the role of user consent have, to date, done little but undermine the usability and user experience of the web.

The cookie consent banners that plague the web perhaps best express this problem. Usually served by little-known middlemen—who operate euphemistically termed 'Consent Management Platforms'—these pop ups make agreeing to widespread tracking as straightforward as possible for users. But the consent that results from the frustrated click of an 'I agree' button, especially when access to content or a service is made conditional on that acceptance, is dubious at best—and presents significant privacy issues[14]. Nonetheless, after consent is obtained on one site, these systems indiscriminately track users across the web, feeding into the system of data brokers previously discussed.

We think fixing this requires a common, transparent design language for user consent. If an organisation wants to make a connection, then users should be actively and freely consenting via an interface that is consistent across organisations, contexts, and operating systems. This cannot be delivered in the web browser, as it often is at the moment, because legitimate consent processes are easily spoofed or emulated by bad actors. By contrast, a universal system would be easy to understand and recognise across the web, reassuring users that a safe and private connection was being made between the organisations they intended.

The final benefit of mandating explicit consent as a point of principle is that it would help ensure that data brokers and other bad actors could not nefariously connect our information up behind the scenes. Because each user would possess a different unique identifier across every relationship in their life, interpolating their actions across services would be exceptionally difficult. Companies that wished to perform such functions would therefore need to seek consent if they wanted to track users. And, of course, only certified providers could ask for such consent under this model—likely making any such tracking far less intrusive than it is today, if regulators so wished.

To further illustrate what stands to be gained, we can look at digital advertising which is facing a crisis, stemming from the sector's continued reliance on personal data. Exploring how this proposal could reform digital advertising illustrates how stepping back and rethinking the foundational assumptions of the web could solve deep-seated problems.

DIGITAL ADVERTISING CASE STUDY

Roughly half of every dollar spent placing an advert currently services the complex network of platforms and exchanges that makes up the adtech industry[15]. Yet with 32% of UK internet users now blocking ads due to privacy and usability concerns, it is clear that something needs to change[16]. The sector is consequently feeling the pressure, particularly as regulators around the world look to strengthen their

[14] https://www.privacyinternational.org/explainer/2975/most-cookie-banners-are-annoying-and-deceptive-not-consent
[15] https://www.isba.org.uk/system/files/media/documents/2020-12/executive-summary-programmatic-supply-chain-transparency-study.pdf
[16] https://www.statista.com/statistics/351862/adblocking-usage/

anti-competition and privacy laws.

Much of the trouble today stems from an over-reliance on personal information, bargained over through opaque auctions. The issue manifests most obviously in the ***bidstream***, which releases metadata including a user's IP address, location, and demographic information onto the open ad market. Via advertising exchanges, supply and demand side platforms then use this personal data—together with inferences purchased from data brokers—to target us, auctioning ads to the highest bidder in the milliseconds it takes for a page to load.

Whether users have meaningfully consented to all these actors using their data in this way is highly contestable. Although bidstream data is supposedly highly regulated, unscrupulous ad buyers can easily extract valuable insights or misuse that data to commit programmatic ad fraud. But, even when the system does work as intended, it fundamentally relies upon middlemen collecting, processing, and monetising personal information, mostly without users' consent or awareness. Indeed, we saw earlier how the 'consent management platforms' that provide legal cover for these practices are barely fit for purpose.

The system we propose would therefore bring a far more logical structure to the advertising sector, replacing its fragmentation and opacity with a new model built around clarity, control, and privacy. As always, this would start by supplanting the exchange of personal information with trusted connections. Instead of tracking users via cookies and IP addresses, direct connections would be established between advertisers and individual users—inculcating a radically different model to that which currently underpins the advertising stack.

The result would look something like this. If a website wished to be ad-supported, the organisation running it would need to be certified to request ads. The first time a user visited a site, they could consent to seeing ads via the same OS-level pop-up that underpins all connections[17]. This screen would be pre-populated with a list of all the organisations that already had a relationship with the user and wished to advertise to them. Users would accordingly be able to deselect any organisations they did not wish to hear from, removing them as a potential ad provider.

As a user browsed the web, their device would then match advertising supply to demand, essentially acting as an advertising exchange. This would cut out most of the middlemen currently involved in advertising, to the benefits of advertisers, ad supported sites, and users' privacy. Most importantly, it would become next-to impossible to track a user across the web without first seeking their permission. The use of unique identifiers would prevent middlemen from profiling users, reasserting the importance of consent as the basis for legitimate online interactions and information processing.

Websites would fill ad space in two ways. The first would involve organisations leveraging relationships they already had to serve ads directly to their customers across the web. A company that you shop with, for instance, could bid to place their ads on other websites you visited. This would be hugely attractive for many organisations. By removing the intermediaries that currently come between them and their customers, advertisers would regain control of those relationships and avoid brand dilution. They would also know their ads were being shown to users that were actually interested in seeing them, which is likely to increase conversion rates. All this means their ad spend would go further—not least as they would no longer be paying for poorly-targeted ads, delivered by a complex network of middlemen[18].

Not all organisations that wished to advertise would already have a trusted connection with potential customers, though. This is where the second option would come in. Entities that already had relationships with many users would be well-placed to run what amounts to a new kind of demand side platform—but one that only operated with the explicit consent of users. Social media sites and search engines, for instance, have direct relationships with large numbers of users. So long as they were certified to show ads, these organisations could therefore leverage these relationships to place relevant ads on behalf of other organisations.

This would replace the open marketplaces for user data that exist today with a closed and direct, consent-based system. There would be no brokering of personal information by third parties. If a company wanted to act as an ad platform, all the ads they placed would therefore need to be branded with the platform's name. This would make it clear to users which relationships in their life were underpinning the particular ads they are seeing.

As a result, data sharing within the advertising ecosystem would be logically perceivable by users. This should incentivise organisations to better respect users and keep the quality of the ads they

[17] This would most likely grant consent for advertising to be shown for a period of time, say three months, before consent would need to be sought again.

[18] Furthermore, direct connections would introduce the possibility of building in better measures for checking an ad has actually been seen, to better target ads and help cut down on ad fraud. Such reporting measures would bring more transparency to advertising, for both users and advertisers.

were hosting high. After all, if a provider consistently showed poor-quality, invasive, or creepy ads, users would likely withdraw their consent and stop seeing ads from that platform. We would therefore expect users to gradually regain faith in the advertising system and begin to exercise newfound control over the ads they saw, knowing that their personal data was no longer being exchanged behind the scenes. The result would be a highly flexible system that put people back in charge of their relationship to advertisers[19].

These points alone would amount to a huge leap forward over the existing advertising model. For organisations that wished to be ad-supported, it would become far easier to integrate legitimate, trustworthy and relevant ads from many providers into their sites via standardised APIs. For the same reasons, it would also become far easier for organisations that wanted to provide ads to place them on a wide range of websites. Both kinds of organisations would therefore prevent a staggering amount of their ad spend from being wasted on middlemen—boosting revenues while cutting costs. Indeed, with an average of 49% of advertiser spend currently going to this network of adtech intermediaries, the savings on a typical advertising budget would be enough to double the number of adverts shown during a campaign[20].

What is more, both ad providers and ad-supported sites would likely find users more likely to engage with their ads in the first place. Via the certification process, regulators would also gain new tools for reigning in excessively privacy-infringing or otherwise undesirable advertising practices. Users would know that the ads they were seeing were coming from certified organisations, creating a better advertising model and a better web.

[19] Of course, this proposal is not just about advertising. It is worth noting that the same API would allow users that did not want to see ads at all to pay for ad-free access, perhaps via microtransactions, should organisations wish to offer this option.
[20] https://www.isba.org.uk/system/files/media/documents/2020-12/executive-summary-programmatic-supply-chain-transparency-study.pdf

CONCLUSION

The notion of a ***common carrier*** has existed in British and American law for well over a century, and ensures that any business transporting something—be it goods, people, or data—must do so agnostically. That is, such businesses may not discriminate; they have a legal duty to carry any lawful cargo, at a common, fair price. A railway operator, for instance, could not legitimately refuse to transport deliveries for a competitor.

In many ways, our devices have become the thoroughfares of our digital lives. So far, though, we have allowed the OS providers to capitalise on their position, often at our expense. The iPhone's Near Field Communication reader, for example, can only be used by Apple's own payments service, preventing users from choosing alternative, competing payment options[21].

Common carrier laws, which already apply to internet and communications companies, provide clear precedent here[22]. If extended to these proposals, such laws would ensure that device manufacturers and OS providers could not discriminate or limit the connections that users wished to make, or dictate how options were presented. Rather than favouring their own services, they would be legally required to treat all possible options equally, putting users back in control.

While these companies could of course be compelled, through legislation, to route standardised requests between certified organisations, we believe that it would in fact be in their economic interest to do so. The early days of the iPhone present a pertinent analogy here. To begin with, the iPhone operated as a closed system, with only a few purpose-built apps and no App Store. Yet Apple soon realised that allowing developers to create apps was a better long-term strategy, even though most would earn them little-to-no commission. This soon led to the creation of innovative new products and services that caused the utility and perceived value of the iPhone to skyrocket, cementing its historical importance.

The system we propose can be expected to increase user privacy, ensure greater interoperability, and require meaningful consent online. Many of these benefits would flow directly from incorporating and building upon existing best practices in cybersecurity competition policy, and interface design—we are not reinventing the wheel. But these practices have not, until now, been integrated into a proposal that combines such technical advances with the wider governance and regulatory structures needed to fundamentally reimagine how personal information works on the web. This proposal takes this latter step for the first time.

Fifty years after the first internet transmission was sent from Room 304 at the University of California to the Stanford Research Institute, the web has become a core part of our social infrastructure, embedded into the fabric of our daily lives. If we are to preserve the founding values of an open, interoperable web in the face of such profound change, we must update the institutions, regulatory regimes, and technologies that make up this network of networks. As we have shown, many of the problems we face stem from the vast amounts of personal information that currently flow through the internet—and fixing this fundamental flaw would have a profound effect on the quality of our lives and the workings of the web.

[21] https://ec.europa.eu/commission/presscorner/detail/en/ip_22_2764
[22] http://www.timwu.org/network_neutrality.html

for or directed toward commercial advantage or private monetary compensation. The exchange of the Work for other copyrighted works by means of digital file sharing or otherwise shall not be considered to be intended for or directed toward commercial advantage or private monetary compensation, provided there is no payment of any monetary compensation in connection with the exchange of copyrighted works.

c If you distribute, publicly display, publicly perform, or publicly digitally perform the Work or any Collective Works, you must keep intact all copyright notices for the Work and give the Original Author credit reasonable to the medium or means You are utilizing by conveying the name (or pseudonym if applicable) of the Original Author if supplied; the title of the Work if supplied. Such credit may be implemented in any reasonable manner; provided, however, that in the case of a Collective Work, at a minimum such credit will appear where any other comparable authorship credit appears and in a manner at least as prominent as such other comparable authorship credit.

5 Representations, Warranties and Disclaimer

a By offering the Work for public release under this Licence, Licensor represents and warrants that, to the best of Licensor's knowledge after reasonable inquiry:

i Licensor has secured all rights in the Work necessary to grant the licence rights hereunder and to permit the lawful exercise of the rights granted hereunder without You having any obligation to pay any royalties, compulsory licence fees, residuals or any other payments;

ii The Work does not infringe the copyright, trademark, publicity rights, common law rights or any other right of any third party or constitute defamation, invasion of privacy or other tortious injury to any third party.

b Except as expressly stated in this licence or otherwise agreed in writing or required by applicable law, the work is licenced on an 'as is' basis, without warranties of any kind, either express or implied including, without limitation, any warranties regarding the contents or accuracy of the work.

6 Limitation on Liability

Except to the extent required by applicable law, and except for damages arising from liability to a third party resulting from breach of the warranties in section 5, in no event will licensor be liable to you on any legal theory for any special, incidental, consequential, punitive or exemplary damages arising out of this licence or the use of the work, even if licensor has been advised of the possibility of such damages.

7 Termination

a This Licence and the rights granted hereunder will terminate automatically upon any breach by You of the terms of this Licence. Individuals or entities who have received Collective Works from You under this Licence, however, will not have their licences terminated provided such individuals or entities remain in full compliance with those licences. Sections 1, 2, 5, 6, 7, and 8 will survive any termination of this Licence.

b Subject to the above terms and conditions, the licence granted here is perpetual (for the duration of the applicable copyright in the Work). Notwithstanding the above, Licensor reserves the right to release the Work under different licence terms or to stop distributing the Work at any time; provided, however that any such election will not serve to withdraw this Licence (or any other licence that has been, or is required to be, granted under the terms of this Licence), and this Licence will continue in full force and effect unless terminated as stated above.

8 Miscellaneous

a Each time You distribute or publicly digitally perform the Work or a Collective Work, Demos offers to the recipient a licence to the Work on the same terms and conditions as the licence granted to You under this Licence.

b If any provision of this Licence is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this Licence, and without further action by the parties to this agreement, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

c No term or provision of this Licence shall be deemed waived and no breach consented to unless such waiver or consent shall be in writing and signed by the party to be charged with such waiver or consent.

d This Licence constitutes the entire agreement between the parties with respect to the Work licenced here. There are no understandings, agreements or representations with respect to the Work not specified here. Licensor shall not be bound by any additional provisions that may appear in any communication from You. This Licence may not be modified without the mutual written agreement of Demos and You.

DEMOS

Demos is a champion of people, ideas and democracy. We bring people together. We bridge divides. We listen and we understand. We are practical about the problems we face, but endlessly optimistic and ambitious about our capacity, together, to overcome them.

At a crossroads in Britain's history, we need ideas for renewal, reconnection and the restoration of hope. Challenges from populism to climate change remain unsolved, and a technological revolution dawns, but the centre of politics has been intellectually paralysed. Demos will change that. We can counter the impossible promises of the political extremes, and challenge despair – by bringing to life an aspirational narrative about the future of Britain that is rooted in the hopes and ambitions of people from across our country.

Demos is an independent, educational charity, registered in England and Wales. (Charity Registration no. 1042046)

Find out more at **www.demos.co.uk**

DEMOS

PUBLISHED BY DEMOS JUNE 2023

15 WHITEHALL, LONDON, SW1A 2DD
T: 020 3878 3955
HELLO@DEMOS.CO.UK
WWW.DEMOS.CO.UK