

Information Architectures: A Framework for Understanding Socio-Technical Systems

Paul E. Smaldino^{1,2,*}, Adam Russell³, Matthew R. Zefferman⁴, Judith Donath⁵, Jacob G. Foster^{2,6}, Douglas Guilbeault⁷, Martin Hilbert⁸, Elizabeth A. Hobson⁹, Kristina Lerman³, Helena Miton⁷, Cody Moser¹, Jana Lasser^{10,11}, Sonja Schmer-Galunder¹², Jacob N. Shapiro¹³, Qiankun Zhong¹⁴, Dan Patt¹⁵

¹ Department of Cognitive and Information Sciences, University of California-Merced

² Santa Fe Institute

³ Information Sciences Institute, University of Southern California

⁴ Department of Defense Analysis, Naval Postgraduate School

⁵ Berkman Klein Center, Harvard University

⁶ Departments of Informatics and Cognitive Science, Indiana University-Bloomington

⁷ Graduate School of Business, Stanford University

⁸ Department of Communication, University of California-Davis

⁹ Department of Biological Sciences, University of Cincinnati

¹⁰ IDea_Lab, University of Graz

¹¹ Complexity Science Hub Vienna

¹² Department of Computer & Information Science & Engineering, University of Florida

¹³ Department of Politics, School of Public and International Affairs, Princeton University

¹⁴ Max Planck Institute for Human Development

¹⁵ Hudson Institute

* Corresponding author. Email: paul.smaldino@gmail.com

Abstract:

A sequence of technological inventions over several centuries has dramatically lowered the cost of producing and distributing information. Because culture and societies ride on a substrate of information, these changes have profoundly impacted how we live, work, and interact with each other. This paper explores the nature of *information architectures* (IAs)—the features that govern how information flows within human populations. IAs include physical and digital infrastructures, norms and institutions, and algorithmic technologies for filtering, producing, and disseminating information. IAs shape everyday lives and cultural practices; they can reinforce societal biases and discriminations and lead to prosocial outcomes as well as unintended social ills. IAs have culturally evolved rapidly with human usage, creating new affordances and new problems for the dynamics of social interaction. We explore the significant societal outcomes instigated by shifts in IAs and call for an enhanced understanding of the social implications of increasing IA complexity, the nature of competition among IAs, and the creation of mechanisms for coordination to guide the beneficial use of IAs. Navigating the challenges IAs pose requires creative experimentation, novel tools, and an understanding of societal norms and institutions. Given that the speed and influence of IAs are accelerating, research in this area is increasingly imperative to ensure that the nature of IAs aligns with a collectively desirable future.

"We are but whirlpools in a river of ever-flowing water. We are not stuff that abides, but patterns that perpetuate themselves. A pattern is a message, and may be transmitted as a message."
—Norbert Wiener (1950) *The Human Use of Human Beings: Cybernetics and Society*

1. Introduction

Technology, human society, and communication networks have always been intertwined—from the earliest footpath networks that enabled trade between small groups of people to today's globalized economy connected by Internet-enabled smartphones. However, the tremendous growth in the scale and capabilities for information exchange has yielded something new that is not yet fully understood. As Philip Anderson (1972) famously observed of physical systems, "more is different." The same is often true of social systems. Building on the sciences of the early 19th century—physics, chemistry, thermodynamics, electromagnetics—humanity built amazing inventions, creating the lightbulb, telephone, automobile, and airplane in a 30-year span. But it wasn't these inventions alone that drove economic growth and transformed society; it was also the networks they spawned—an electrical grid, a telephone network, a highway system, and an air transportation system. With the growth of information technologies, these networks became enmeshed into the fabric of our social and cultural organizations. Most recently, AI algorithms, including large language models (LLMs), have become rapidly adopted across many industries and threaten further disruption of the social fabric through their influence on information flow (Brinkmann et al 2023).

Over the past century, the costs of creating, distributing, and processing information have dramatically decreased. These extreme discounts have disrupted the economics of socio-technical systems, stripping distributors like newspapers, telephone companies, or TV networks of power and fomenting new competitions for human attention. In just the last century, we have gone from a small number of outlets competing for overlapping market share to a large number of outlets exploiting niche audiences to competing repositories where anyone can generate and disseminate content (Wu 2016, DiResta 2024). While an individual in a developed country can refrain from using certain technologies—e.g., abstaining from smartphone use or subscription payments—doing so increasingly means effectively withdrawing from society. Engaging in trade, obtaining housing or education, and even finding romantic partners all depend on some substrate of underlying technical and information networks. The control over the flow of information—which now includes not only direct communication but nearly all modern exchanges of financial and social capital—is one of the pivotal levers involved in the power struggles of contemporary geopolitics (Farrell & Newman 2023).

Technologies are, for the most part, neither good nor bad by themselves, but they do change how people can organize, communicate with, and even exploit one another. As McLuhan remarked, "We shape our tools and thereafter they shape us" (Culkin, 1967). On one side, there are the considerable economic gains that have come with technology. There is much less global poverty, childhood death, and arguably global suffering than there was two centuries ago. On the other hand, technologies have potentially exacerbated a slew of social ills, including political

polarization (Lelkes et al. 2017), socioeconomic inequality (Acemoglu & Restrepo 2018), mental health crises (Scott et al. 2017), and societal discord (DiResta 2024). We are also beginning to recognize that there are competing architectures, both within and between states, with implications for the future of choice with respect to values and ethics.

If human societies and information technologies are inexorably coupled, then it is imperative to understand the dynamics involved in those couplings, and how the specific features of the information technologies that tie a group of people together shape and are shaped by those same people. We use the term *information architecture* (IA) to concisely describe this concept. There is already a great deal of research into techno-social systems, with burgeoning studies of social networks, mesoeconomics, and cultural evolution (Bakshy et al., 2012; Centola 2015; Acerbi 2019; Farrell & Newman 2019; Wagner et al. 2021; Brinkmann et al. 2023). However, there are also many gaps in this work, and few efforts to understand how different information architectures interact, cooperate, and compete. We offer the beginnings of a unifying, interdisciplinary framework to help us build a desperately needed social science for the 21st century.

2. What Are Information Architectures?

The term “information architecture” is not new, and dates back to at least the 1970s; it has been used within some communities of organizational science, urban design, software development, and user experience (UX) research (Brancheau & Wetherbe 1986; Dillon & Turnbull 2005, Morville & Rosenfeld 2006; Downey & Banerjee 2010; Arango 2011). In these contexts, information architecture typically refers to a process of intentional, active design as performed by one or more “information architects.” A representative definition of information architecture in this literature is the “organization of information to support findability, manageability and usefulness from the infrastructural level to the user interface level” (Downey & Banerjee, 2010). The information architect's job, in this context, is to design systems with users in mind, to facilitate the best possible user experience and the highest-value engagement with information in those systems (Arango 2011).

We aim to extend the traditional conception of IAs from this literature in two ways: first, to expand it beyond the scale of specific organizations, applications, or interfaces to encompass the flow of information throughout society, and second, to expand the concept of “design” to include non-deliberate processes. We are interested in information architectures that result from the combination of both deliberate and non-deliberate processes that unfold over multiple time scales. Although our perspective on information architectures differs from past usage, we adopt the term because, at heart, we still seek to describe and understand how the nature of social and technological organization facilitates how information is or is not accessed and transmitted.

For the remainder of this paper, we define an information architecture as *the set of structural features and affordances that enables and constrains information flow between sources and receivers within a population*. By structural features, we include both physical infrastructure—such as transportation networks, broadcast media, telecom and cell networks, geographical

features (e.g., mountains dividing populations, rivers connecting them)—and cultural features, including social norms, laws, languages, and both formal and informal institutions. We also include digital infrastructures and other mechanisms for online communication via computers or smartphones, including peer-to-peer communication channels (e.g., SMS, WhatsApp, WeChat), social media platforms (e.g., Facebook, Instagram, Telegram, Twitter/X, TikTok), blogs and related platforms (e.g., Medium, Substack), and other conduits for assessing or storing information (e.g., blockchain, banks, public records, etc.). Such features can enable information flow (e.g., by allowing individuals to send messages or access stored information) and constrain it (e.g., by restricting certain types of content or the set of individuals who can send or receive messages or access certain information). By affordances, we mean the opportunities and behaviors made possible by those structural features (Gibson 1979; Norman 1990). Given the extent to which our modern world is connected, there is in some sense only one all-encompassing information architecture, encompassing every individual on the planet. However, this perspective ignores the extent to which subpopulations can be functionally defined based on national, linguistic, ethnic, and other boundaries. A focus on IAs may provide insight into the diversification and maintenance of subpopulations. For example, as groups develop different IAs, their organizational capabilities change and they increasingly undergo different evolutionary trajectories.

There is a large literature on the relationship between the structure of social networks and the diffusion of information on those networks (Granovetter 1973; Watts 1999; Young 2006; Zollman 2013; Sáenz-Royo et al. 2015; Centola 2015). Indeed, information architectures include conduits for diffusion, but the characterization of an information architecture goes beyond merely what is being diffused and the extent of that diffusion. Such a characterization also includes consideration of the identity of individuals transmitting or receiving information, the accessibility of information, and how the packaging of information can change its nature. A key aspect of our treatment of information architectures is how they interact with one another, a feature usually omitted from discussions of diffusion. Even when diffusion studies consider multiple groups, they rarely highlight differences in how information flows both within and between groups (for some exceptions, see Smaldino & Jones 2021; Turner et al. 2023; Milzman & Moser 2023). Information architectures can therefore be characterized along several dimensions, and these means of characterization are important for comparing and contrasting distinct architectures and considering how they might compete.

3. Characterizing Information Architectures

As an illustrative model, consider two fictitious societies. These are caricatures but represent two extreme poles of information control, which will help us explore the consequences of the institutions that shape the information architectures within each of them.

The first of these is called *Harmonia*. Its institutions are intended to produce an absence of discord in the interest of social stability—carefully filtering outside information and regulating internal interactions, curation algorithms, and content to weed out misinformation or framings that might sow discontent. Because of the vast complexity of these tasks and the difficulty of

employing consistent means of decision-making, control is centralized with the leadership of Harmonia. Naturally, this leadership will wield its power over the content people can consume with great care, while the interests of the leadership and the people—in terms of concerns for accuracy, well-being, and freedom—will not always align.

The second of these societies is called *Libertaria*. Its institutions value individual choice above all, resisting interventions or regulations by leaders concerning information flows. Because any imposition of control has the potential for abuse and quashing of minority voices, no top-down control is granted to any central authority, though inequality of wealth inevitably leads to inequality of influence. The people in Libertaria are thus immersed in a sea of misinformation, disinformation, and attention-maximizing algorithms, but they are also free to speak their minds and to seek out those whose beliefs and goals align with their own, even when these are at odds with societal harmony.

As exemplified by our examples of Harmonia and Libertaria, the study of information architectures concerns itself with understanding the relative benefits and costs of living in a particular society, the interactions between societies, the interventions that citizens and leaders might engage in, and the ability to transition between alternative structures to address functional needs or respond to shocks (e.g., during wartime, Libertaria may feel compelled to adopt some of the strategies of Harmonia, while the pursuit of scientific progress may compel Harmonia to learn from Libertaria). It might be that certain kinds of IAs enable greater flexibility and reversibility of architectural features than others. Information architectures can reinforce societal biases and discriminations and lead both to prosocial outcomes and unintended social ills. Because information architectures are so tightly linked to economic activity, regulations that quell one social ill in the near term may have an inadvertent impact on social gains in the long term. A society may burn books to quell discord in the short run, only to discover that the information in those books was necessary to solve a future problem that has more enduring consequences.

Information architectures surely vary along many dimensions, and useful characterization of information architectures should allow us to highlight their differences. Here we highlight four dimensions along which we might characterize IAs, with the understanding that this list is not exhaustive. Indeed, future work should more precisely capture the “computational primitives” (cf. Coward 2014) that allow elements of information architectures to create, redirect, organize, filter, curate, modify, or destroy information.

Flows and structures. One way to characterize IAs is by their permissible and actual paths for information flow. Who is connected to whom? What are the possible paths for communication between individuals A and B? Flow entails directionality and often hierarchy. For example, broadcast media is characterized by one-to-many information flows, whereas peer-to-peer networks are many-to-many. This has changed, as community opinion leaders changed from well-informed individuals who consume news to well-connected individuals who create engaging content on social media (Wu 2016; DiResta 2024). As with many complex systems, idealized formal models may help to characterize information architectures. Formal structures such as

multiplex, multipartite networks may be useful in this regard, with nodes representing individuals, organizations, and other information sources and edges representing connections via various communication channels, including physical proximity (Atkisson et al. 2020). In the case of our fictitious societies, Harmonia's information architecture would feature centralized and regulated information flows controlled by leaders while Libertaria would be characterized by a decentralized structure with open, peer-to-peer communication channels.

Volume and bandwidth. How much information can be stored within various media and transmitted between individuals, and at what rate does that transmission occur? In the past, information was severely limited by spatial proximity and by human physical and perceptual capacities, despite the importance of technological advances like the printing press (Jara-Figueroa et al. 2019). Now, beyond the changes brought about by advances in transportation and physical infrastructure, the internet and digital technologies allow storage and rapid communication of nearly limitless amounts of information across nearly any physical distance. In the case of Harmonia, we might find limits to the volume and bandwidth of information transmitted imposed to emphasize circulation of approved content during monitorable hours, in contrast to Libertaria's unrestricted approach allowing for information sharing up to the limits of supporting technical networks. It bears mentioning that the volume and bandwidth of IAs are also set by physical limits in the amount of energy required for an IA to gather, process, store, and distribute information. This is highlighted by current debates surrounding the ecological footprint of cryptocurrency and large language models (Luccioni et al., 2023; Zribi et al., 2023). In this way, both ecological and sociological factors also define the volume and bandwidth constraints of an IA.

Distributional constraints. Not all information can be shared with the same ease. There are often legal and normative constraints over words, images, topics, and sentiments whose transmission are either encouraged or forbidden in certain contexts or by certain classes of people. In Germany, for example, imagery and sentiments associated with the Nazi regime are proscribed, and most direct advertisement of medical products is prohibited throughout the European Union. In the US, citizens can say nearly anything they want about politics but are restricted from making false claims about medical products, while non-citizens have limited rights regarding political speech. The Chinese government censors references to content, including the Tiananmen Square protests and massacre, Tibet and Hong Kong independence movements, state persecution of Uyghurs, and the resemblance between Chinese President Xi and Winnie the Pooh. Social media platforms may restrict or hide posts flagged as containing inflammatory content (Gillespie 2018; Gorwa 2024). Logistical constraints can also influence distribution. For example, during the COVID pandemic, movies could still be released online while museums and the performing arts suffered from severely decreased patronage. Thus, IAs can be characterized by how different types of information content can be distributed between sets of nodes. Harmonia would be characterized by strict constraints on content distribution, filtering out potentially disruptive information, whereas Libertaria would allow all types of information, including potential misinformation, to circulate freely.

Accessibility. Consider the idealized case in which all the information an individual could hypothetically access in a society is represented as a massively long bitstring, with each bit representing a piece of information. An individual's accessibility is the proportion of that bitstring that is available to them. This thought experiment allows us to consider how accessibility varies across the society, and to characterize inequality in access to information as a function of characteristics such as wealth, ethnicity, or social class. Accessibility may include access to archives or repositories, including libraries, websites, website APIs, code repositories, and digital memory banks—we might consider access to previously stored information as the *retrievability* of the IA. Accessibility can change over time for individuals as well as societies, both because new information is created and old information is destroyed, and also because infrastructures, laws, and norms may change. In *Harmonia*, information access is tiered and controlled based on role, while *Libertaria* theoretically offers open accessibility to all information, though practical access may be influenced by individual resources and private monopolization.

Our examples of *Harmonia* and *Libertaria* are not meant to suggest that certain IA features are compatible only with democratic or authoritarian states. For example, a state-controlled authoritarian environment may co-opt an IA that has a similar network structure, bandwidth, and accessibility to that of *Libertaria*-like society. It may adopt a “firehouse of falsehood” strategy to “flood the zone” of a relatively open IA with misinformation to bury the truth in a flood of lies (Paul & Matthews 2016, Zegart 2022). The intended result is widespread distrust of any information. Indeed, Hannah Arendt (1951) suggested that this type of dynamic may underpin the rise of totalitarianism from more democratic societies.

How and why do IAs with various characteristics come about? How do they shape human interaction, and how are they in turn shaped by those interactions? We turn to these questions in the next two sections. We will first use the characteristics described above to explore historical cases in which new information architectures led to wider downstream consequences. Then we will turn to elements of more contemporary information architectures, with a focus on internet applications.

4. IA Examples and Illustrative Case Studies

In order to highlight the importance of understanding the dynamics of IAs, this section provides a series of historical comparisons within IAs at major moments of change and across IAs. We will look at shifts induced by the introduction of the first transatlantic cable in 1866 and by the introduction of cellular phone coverage along the coast of India from 1997-2001 and in Iraq from 2003-2008 during the civil war. For comparisons between contemporaneous societies with different IAs, we will look at differences in scientific progress between England and Spain in the 17th century.

The laying of the first successful transatlantic cable in 1866 marked a clear shift in the information architecture for the cotton trade. As discussed by Steinwender (2018), the cable reduced the time for information transmission from 7-15 days to just one day. In terms of the four dimensions outlined in Section 3, the immediate effects were most obviously seen in

changes to bandwidth, but the introduction of rapid transatlantic communication affected all four factors and created new affordances for further changes to come:

1. Flows and structures: The middlemen of shipping and postal carriers were shortcutted with the facilitation of more direct communication channels. Nevertheless, connectivity between principal actors in the cotton trade was not different in shape.
2. Volume/bandwidth: There was a massive increase in the transmission rate for low data volume communications, e.g. short telegrams regarding prices or financial conditions. The ability to transmit large amounts of data (e.g. documents) was unchanged.
3. Distribution constraints: Industries with early access to cable transmissions gained an advantage, while the reduced economic costs of communication opened the door to greater competition.
4. Accessibility: The transatlantic cable granted the connected actors (and nations) increased access to each others' information, increasing the inequality in information access.

This new architecture had dramatic effects on economic flows and structures (Steinwender 2018). Price dispersion between New York and Liverpool markets dropped substantially—the average price difference fell by 35%, the equivalent of removing a 7% ad valorem tariff. Trade volumes also increased by 37% on average after the cable was introduced. The architecture enabled more efficient alignment of supply and demand across continents, with estimated efficiency gains of 8% of annual cotton export value from the US to Britain.

The information architecture enabled by undersea cables catalyzed major changes across multiple domains beyond economics. For example, they transformed business operations and logistics. News reporting also drastically changed, with outlets like Reuters even investing in cables to capitalize on access to faster international news (Headrick and Griset 2001). More broadly, it impacted politics, culture, and society. The ability to almost instantaneously transmit messages revolutionized communications between governments, businesses, and individuals.

Over a century later, the introduction of cellular phone coverage around the world had similar impacts by enabling new kinds of immediate information flows. In terms of our four dimensions, cell phone coverage introduces a broader set of changes:

1. Flows and structures: Massive increase in interconnectivity.
2. Volume/bandwidth: Massive increase in rate for even high-bandwidth data communications.
3. Distribution constraints: Large increase in the ease of communication with anyone whose phone number is known, reducing reliance on face-to-face interaction.
4. Accessibility: Cellular customers could access information in more contexts and locations, while costs of cellular service introduced inequalities in communication access.

Not surprisingly, the adoption of mobile phones catalyzed major shifts in market dynamics and economic outcomes. One study showed that when cellular coverage was introduced along the coast of Kerala, India, from 1997-2001 the price dispersion between local fish markets dropped

substantially, with the coefficient of variation in prices falling from 60-70% to 15% or less almost immediately (Jensen 2007). This enhanced integration enabled fishermen to eliminate the previous 5-8% of daily catch that was previously wasted due to saturation in local markets. Consumer prices for fish declined by 4% on average, while the improved efficiency increased fishermen's profits by 8% on average, with those using phones gaining the most. Some fishermen reported compiling extensive contact lists with hundreds of potential buyers to survey prices before deciding where to sell their catch. Similar effects have been observed in other developing country agricultural markets (Aker 2010).

The ability to coordinate sales remotely revolutionized trading, transforming isolated village fish economies into integrated networks. Beyond the economic shifts, mobile connectivity impacted social relations and status. Fishermen with mobile phones rose in prominence, with their broader set of contacts making them valuable nodes for transmission of market information, even to those without phones. The examples of the telegraph cable and adoption of mobile phones in Kerala demonstrate vividly how new architectures that enhance information flows can profoundly reshape economic and social landscapes.

In Iraq from 2003-2008, the gradual rollout of cellular coverage had more complex effects. There, insurgents used cell phones to “coordinate actions, execute attacks, and quickly react to counterinsurgency operations” and even threatened providers for not doing enough to maintain their networks (Shapiro & Weidmann 2015). From the perspective of Iraqi government forces and their U.S. allies, cellular coverage also made “it easier for the population to share information about insurgent activity, and to safely and anonymously call in tips” on insurgent activity (Shapiro & Weidmann 2015). The shifted IA created by cellular coverage offered select advantages to both sides, with fewer attacks overall and some evidence that insurgents shifted to using more indirect fire weapons, which could be used to attack from locations of the insurgent’s choosing.

Turning to comparisons across countries, it is instructive to compare England with Spain in the 1700s. In terms of our four dimensions, there were some clear differences, but also many similarities:

1. Flows and structures: Both nations had similar socioeconomic and technological constraints, though differences in culture and religion created somewhat different patterns of information flow.
2. Volume/bandwidth: Both nations had access to similar technologies and affordances for the storage and transmission of information.
3. Distribution constraints: Many ideas faced significant content-based restraints in Spain due to the ongoing influence of the Inquisition.
4. Accessibility: Spain had substantially more universities (Ruegg 1996) but lacked the government-funded scientific societies and social norms that made engagement in science a proud hobby for many wealthy individuals in England (Shapiro 1971).

If one had to guess in 1650 which country—Spain or England—would lead the flowering of science, choosing Spain would have been entirely reasonable. The country had a larger population, roughly 10 times as many universities, and a similar GDP per capita adjusting for

purchasing power. England, however, established a nationally supported scientific society in 1660 (Merton 1970), putting the state's imprimatur on the notion of open scientific discourse. From 1660 through 1700 Spain remained in the grip of the Inquisition, and did not establish a national scientific society until 1790 (De Vos 2006). The subsequent difference in these countries' contributions to science was massive. Seventeenth-century England saw a flourishing of research with its scientists playing leading roles in many fields, while Spain experienced a century of stagnation and declining GDP, with the Inquisition leaving persistently lower levels of education and social trust in its wake (Drelichman, Vidal-Robert, and Voth 2021).

5. Feedback Between Information Architectures and Individuals

Information architectures structure and influence the social activities of the actors embedded within them in at least two distinct ways. First, IAs enable social contact by providing channels through which information can flow. Second, IAs store information, granting actors powers of memory and recall beyond their otherwise immediate surroundings. When combined with IAs' ability to scale communication, this often allows for cheap—albeit biased—information search. This in turn can influence the establishment and execution of social and political institutions, shaping peoples' everyday cultural practices and evolving as the technologies, institutions, and practices that enable information architectures evolve.

Modern-day internet applications are important facilitators when it comes to the fulfillment of human needs to communicate, form relations, cooperate, and coordinate. For example, as of 2017, dating apps have become the primary means for heterosexual people in the U.S. to find romantic partners (Rosenfeld et al. 2019). While some people do still find partners “the old fashioned way” (e.g., via shared activities or mutual friends), dating apps have taken over a central role in fulfilling the basic human need of finding a partner, likely because they reduce the costs of both search and communication (Sumter & Vandenbosch 2019). Because online dating lowers the cost of finding new contacts, it also affords a cultural practice of “ghosting,” which was relatively rare in the pre-app era (LeFebvre 2019, Timmermans et al. 2020). The restructuring of information architectures has led to similar outcomes in other times and places. Among the Magar people of Nepal, for example, the introduction of mail and postage in their valley territory in the 1980s and '90s restructured their system of arranged marriages, leading to the increase in elopements as a result of more efficient search among potential spouses (Ahearn, 2001). Online, self-organized communities on Reddit are established through efficient information exchange (Krohn & Weninger, 2022), providing support for topics and norms that would not typically be afforded in non-urban, geographically-isolated communities in the pre-Internet era.

Online platforms such as LinkedIn, Google Scholar, and ORCID store publicly-available information about individuals' professional personae, making it possible for users to rapidly build and disseminate professional reputations (Lampel & Bhalla 2007, Antheunis et al. 2013, Kim & Cha 2017), an important part of relationship-building that historically took much longer to grow (Hruschka 2010). An even more pervasive example of the entanglement of information architectures in everyday life is the market: prices of goods and services in markets encode

information about their means of production and people's beliefs about their value (Thicke, 2018). Markets organize a broad range of human activities and provide the basis for decision making from acquiring food and energy to labor markets. The information encoded in markets can even be used to predict events such as elections with some success (Hanson, 2007; Sethi et al. 2021), a phenomenon made possible by the rapid aggregation and transmission of market information afforded by modern online technologies.

The shift of much human interaction from ephemeral face-to-face conversation to persistent IA-mediated discussion has brought about ubiquitous “data shadows”—vast archives of one's comments, transactions, and other recordable behaviors. Some of this material, such as saved emails, is managed by the individual; much more of it is stored, owned, and potentially sold by the platforms on which the interaction took place. This has profound implications psychologically, socially and legally as the past, instead of fading into a dim temporal distance, remains ever-present and easily searchable. Statements made for a particular situation and audience can resurface at any time, and in quite different contexts; this immediately changes people's ability to keep information private and, eventually, transforms cultural conceptions of privacy and public space (Spiekerman et al 2015; Donath 2013, ch 11). Thus far, third parties have been the primary users of this information: governments pursuing criminals (or repressing dissent), advertisers targeting consumers; these uses are often harmful to the subject and creators of the information. Potentially, these archives could also be used by and for the benefit of their subjects and creators, e.g for self-reflection, social navigation, or enriching online self-presentation (Viégas et al 2006; Donath 2013).

By aligning information and social practices, the components of information architectures shape how people build their cultural repertoires, networks, and shifting norms and values. In many ways, IAs not only reduce the cost of search, but also bias the transmission of information (Bozdog, 2013; Wu 2016; DiResta 2024). As IAs rapidly transmit information, they tend to amplify some information contents and dampen others; information which is transmitted can therefore reflect biases of the underlying user base—or of the governmental or corporate organizations that control the information flows. This point was made by Herman and Chomsky (1988) in their analysis of the media—multiple stages of filtering produce media employees that promote the interests of the ruling classes without requiring their deliberate intent to do so. More recent analyses have shown that continued increases in the rate of information transmission likely reinforce and amplify such biases (Nikolov et al. 2019). This amplification has been observed in social media newsfeed algorithms (Munn, 2020) and recommendation algorithms for video websites such as YouTube (Ribeiro et al. 2020). In recent years, the amplification of these biases has also been implicated in the spread of misinformation (Srba et al. 2023) and increasing polarization in the United States (Feezell et al. 2021; González-Bailón et al. 2023).

Feedback can lead to co-evolutionary dynamics between IAs and their users, resulting in a shift between the stated purpose of the IA and the ways it is actually used in the real world (Sperber 1994). Because IAs determine how information is transmitted, cultural selection will tend to provide content best suited to the nature of the architectures, and that content may then provide further selection on the nature of the architectures (Centola & Baronchelli 2015; Acerbi 2019; Guilbeault et al. 2021) . While the medium may not exactly *be* the message (*contra* McLuhan

1964), the medium shapes the message and the message shapes the medium. For example, the nature of a market shapes the products exchanged on it while also changing to accommodate new forms of exchange (Boettke et al 2004). In other cases, an initial form can come to be repurposed for other functions, as in the case of hashtags, which originally served a now-defunct role to facilitate information searches in large-scale databases, but are now utilized primarily for linguistic emphasis (Heyd & Puschmann 2017).

The design of information architectures is also rarely value-neutral (Scott, 1998). They are at least partly designed by stakeholders in accordance with their values, goals, and designed outcomes. The resulting structure of an information architecture is therefore the outcome of a dynamic process between the information architecture, those who utilize it, and those who contribute to its design (these last two need not be fully separable). This co-evolutionary dynamic creates at least three points for conflict: tensions between designers and users, between designers and IAs, and between IAs and users. Note that the following examples focus in large part on IA components that were deliberately designed, though it is important to note that much of IA design results from non-deliberate processes that may be unconscious or emergent (Smaldino 2014).

In the first case, designers of IA components may come to find that their product is used in a manner inconsistent with their original intentions. Such cases include the development of Bitcoin. Designed originally to serve as a competing currency with modern fiat currencies, it is now largely utilized as an investment tool rather than as a form of liquid currency (Brunton 2019). Tensions between designers and IA components may result when the structure of the system makes adaptation to fulfill a particular demand difficult (that is, meeting a critical goal is not among the set of what Kauffman (2000) has dubbed “adjacent possible” solutions). Such tensions arose in the relationship between Twitter and its designer Jack Dorsey, who spoke on Twitter’s entrenched design features in 2019, stating, “I don’t think I would even create ‘like’ in the first place because it doesn’t actually push what we believe now to be the most important thing, which is healthy contribution back to the network” (Quito 2019). Twitter is an especially interesting example, given the extreme changes to the social media platform following its acquisition by Elon Musk in 2022 and its subsequent rebranding as “X.” The rules and algorithms changed rapidly, followed by a dramatic decrease in the number of overall users (Ingram 2024), an increase in hate speech propagated on the platform (Hickey et al. 2023), and a demographic shift rightward (McClain et al. 2024).

Feedback between IAs, users, and designers can be cooperative as well as antagonistic. In some cases, features may be added that help the IA propagate information and similarly add features that users requested. On the other hand, unwanted changes can come about which are required by the platforms, but not by users. 2-Factor Authentication, for example, is often perceived by users as tedious, but arises due to misuse of the architecture by a subset of the population. Finally, continued co-evolution can lead to stark changes in the value and affordances of some features. Consider, for example, the design of website “cookies,” which were originally designed to facilitate faster interactions between websites and their users. Over time, these were exploited by advertisers, who used cookies across websites to track user behavior. In response to this, EU legislation arose requiring users to actively grant consent on

websites that used them (Jones 2020). More generally, co-evolution in under-regulated capitalist markets has occasionally (though perhaps increasingly often) led platforms that were initially optimized for user experience to increasingly subordinate that experience in favor of stockholder profits and corporate control, a process which Cory Doctorow (2023) has dubbed “enshittification.”

Another example is the feedback between social media companies, users, content algorithms, and “influencers” (DiResta 2024). Social media companies create content algorithms to increase (or maximize) “user engagement,” which is defined by how much users use the platform. Content that increases engagement also increases revenue for the companies, which share their revenue with the users who generate the most engaging content: the “influencers.” Influencers are incentivized to create content that generates the most engagement because it results in more revenue for them. Influencers who do not continue to produce engaging content will lose influence and revenue. The content algorithms also typically promote content that is getting the most engagement, leading to a “rich-gets-richer” effect with some influencers gaining huge followings relative to the typical user. This effect is observed even in academia: those who share their work on social media and have more followers are more likely to be cited (Chan et al. 2023). This has the potential for feedback in which the existing IAs shape the ongoing research related to the IAs themselves, thereby influencing the future development of IA structures or interventions.

The speed at which information architectures are implemented appears increasingly to be much faster than the speed with which human societies can adaptively respond to them. Systems that were once opt-in can scale rapidly, becoming requirements for human integration into social networks. For example, in many industrialized nations, not having a smartphone creates numerous impediments, some of which have nothing to do with the technology’s traditional functions for person-to-person communication. Increases in scale similarly expose additional marginal downsides. This is exacerbated when control over information architectures is less democratic and egalitarian, and individuals feel less sense of agency over decisions affecting information architectural design. While designers and users are both, at some level, active participants in information architectures and therefore both contribute feedback to their design, participants may feel helpless in terms of their ability to influence IAs or societal dependence upon them.

6. The Nature of IA Competition

Information architectures are not monolithic or uniform. As described above, different societies (or sectors of a society) may be governed by different architectures. As such, IAs may be subject to intergroup competition, much as groups with different norms and institutions have competed throughout human history (Henrich 2004; Smaldino 2014; Richerson et al. 2016). What do we mean when we say that information architectures compete? What do they compete over? What are the outcomes of the competition?

We find it useful to approach these questions from a neo-Darwinian perspective. That is, we conceptualize competition as the product of variation and selective retention. IAs vary over

aspects of their design—for example over their structures, bandwidth, distributional constraints, and accessibility (see Section 2). Some IAs are more likely to persist and expand, and others are more likely to wither or die.

A possible currency over which IAs compete is the flow of information between sources and receivers. Successful IAs are those that capture larger shares of information flows in particular populations or domains. The boundaries of an IA are therefore determined by the boundaries of the information flows. These boundaries need not be completely nonporous, but IAs can be defined as routes by which information flows relatively freely.

As with biological organisms, the evolutionary success of information architectures depends on the environments in which they compete. These environments include not just physical constraints, but the laws, rules, and norms that govern their use—what has been called the “up-code” by Shapiro (2023; in contrast with “down-code”: the operating systems and network protocols that directly shape information flow). For example, in a state-controlled environment, successful IAs would have structural features and affordances that most closely align with the goals of the state (Scott 1998; Bradford 2023). As we described in Section 3, a state-controlled authoritarian environment may, for example, favor one-to-many broadcast IAs where the state can distribute its message to the masses and control the flow of information. The state may also suppress many-to-many IAs, where any person’s messages can spread to any other, even those critical of the authoritarian regime. State control of broadcast media could also be combined with monitoring and suppression of many-to-many IAs, provoking the development of clandestine, though inefficient, IAs that evade government censors (such as VPNs).

In a more capitalist environment, successful IAs might be those that best attract, and especially monetize, flows from consumers for profit (Wu 2016, Bradford 2023). The two main models of this are subscription IAs, where individuals pay to receive and/or transmit information and advertising IAs where the information transmitted is incidental relative to the information that can be captured about users and sold to those who would pay for it. For subscription IAs, those that attract the most subscribers who are able and willing to spend the most money will tend to win. For advertising IAs, those that attract the most attention from the most people willing to spend money on advertised products will tend to win. This is not without its dangers. For example, the need to increase usage traffic may favor algorithmic IAs that attract and maintain attention by provoking outrage in users, which may in turn increase societal polarization when outrage is directed toward members of other social or political groups (Brady et al. 2023; Steiglechner et al. 2023; Lerman et al. 2024). IAs implementing these have been successful competitors against other IAs, but arguably at an overall cost to societal well-being. Additionally, competition over which IAs can best collect and sell personal data will tend to favor IAs that feature the weakest privacy protections possible without provoking backlash from users.

Because senders and receivers must all use the same IAs to transmit information, “networking effects” can create additional inefficient outcomes in IA competition. Networking effects, a form of lock-in, create social inefficiencies because it is difficult for a small number of users to switch to another IA unilaterally. The classic example is the QWERTY keyboard, which was designed to be inefficient, reducing the number of sequential keystrokes close to one another to avoid

jamming on mechanical typewriters. Because most people learn to type with a QWERTY keyboard and it is costly to manufacture and distribute alternate keyboards layouts, this inefficient layout persists well into the digital age. Similarly, large social media platforms are very difficult to unseat by objectively better alternatives because this would require a critical mass of users to make the switch simultaneously. Proposals to legislate interoperability between platforms—making communication between users of different social media platforms as easy as between users of different telephone carriers—would change the consequences of these networking effects and allow smaller competitors to persist more easily (Aral 2020; Doctorow 2020).

Environments less constrained by state control or capitalistic drive might favor IAs that produce public goods and are maintained by donations, volunteer labor, and institutions that favor collective action (Hess & Ostrom 2007). However, there is always the danger that such IAs will succumb to “enclosure” by corporations or states, which often have strong incentives to control information for their own ends (Hess & Ostrom 2003). It has been argued that the European Union has structured a regulatory environment based on human rights rather than either authoritarian or market forces, and that this regulatory regime might ultimately become the primary environment in which IAs compete (Bradford 2020). However, global IAs must compete—just as organisms do—by succeeding in *many* different regulatory environments that impose often-contradictory constraints. For example, backers of a generalist IA might eschew regulation and user privacy protections in a capitalist environment, accept regulation and user privacy in a rights-based environment, and help suppress speech and user privacy in an authoritarian environment.

So far, we have looked at competitions between IAs within specific environments, but as many IAs span the globe, they are increasingly becoming competitive domains between international actors, both state and non-state. Consider competition between states. State institutions set constraints on information architectures within those states. Some states assume control of the information within the state, some cede control largely to corporations. Some have strict privacy controls, others do not.

With the “information environment” becoming a prominent domain of international competition, and with states seeking greater influence over the information flows within other states (Farrell & Newman 2023), it is difficult to predict which IAs will dominate. Will isolationist IAs in North Korea and authoritarian IAs in China maintain state power by limiting the flow of available information from the outside world and between citizens of those states? Will the relatively open IAs of Western countries continue to be susceptible to misinformation campaigns from abroad? Will states with privacy regimes that allow private firms to scoop up personal data be at a disadvantage when this data is sold to competing states for a small cost? When IA competition becomes a proxy for state-competition, selection happens on up-code as well as down-code (Shapiro 2023).

What are the potential outcomes of IA competition? One possible outcome is efficiency; IAs will become better at transmitting, compressing, or filtering information. For example, the printing press, the mail system, the telegram, the typewriter, the telephone network, radio, broadcast

television, cable television, and satellite communications—all of these captured information flows from preceding technologies and allowed for more and more rapid and low-cost communication. However, competition also drives IAs to become more efficient at the socially negative aspects of their objective functions. Outside of contravening forces, state-controlled IAs will become better tools of state control. IAs in capitalistic environments might get more efficient at attracting attention and advertising dollars by serving more and more polarizing content and selling more and more user information.

7. Open Questions and Needed Research on IAs

From the perspective of dual inheritance theory (Henrich & McElreath 2007; Russell & Muthukrishna 2021), humans and IAs have been co-evolving for a long time. This makes sense given that IAs shape human communication, which in turn is entangled with our ability to cooperate and coordinate at scale, to learn socially, and to build (and treat as real) collective fictions like identity, stories, religion, currency, and laws that then act as causes of their own in the world (Henrich 2015; Brand et al. 2021; Newson & Richerson 2021). From this perspective, information architectures are critical for enabling and constraining almost all of our truly social interactions and exchange, from how we coordinate to whom we learn from to which stories or fictions are disseminated and come to dominate populations and systems. Taken to a logical extreme, IAs are not just cultural and societal artifacts but are *intrinsic* components of cultures and societies, since without sufficient information transmission, incorporation, and storage, culture itself can begin to erode (Smaldino & Richerson 2013).

Today's IAs do not seem to be merely more of the same thing we've always had. That said, more of the same thing can be sufficient to provoke the emergence of a *new* thing. The increasing scale, reduced cost of information transmission, growing interconnection, and pervasive reach of modern IAs remind us that “more is different” (Anderson 1972). We need new research on IAs precisely because we cannot depend on historical case studies and precedent to help us understand—much less predict or prescribe for—today's IAs. Their growing interconnectedness and dynamism means we are living in a world shaped by IAs that reflect what Weaver (1948) prophetically called “organized complexity.” Far different from systems of “organized simplicity” (which can be tackled by studying a few variables) or “disorganized complexity” (which avail themselves to statistical methods), systems characterized by organized complexity, with its numerous variables “interrelated into an organic whole”, require new methods and approaches that include the use of mathematical models, computer simulations, and, perhaps most importantly, interdisciplinary collaboration among experts from different fields (March & Simon 1958).

We are essentially trying to build a science to better understand IAs that are evolving at the speed of engineering—a speed that historically far outstrips the pace of science. This engineering is ushering in new IAs that are bringing humans and increasingly capable machines together to create new kinds of information storage, computation, and communication processes and capabilities (Brinkmann et al. 2023). Indeed, we may be witnessing the birth of new kinds of IAs. Further, IAs' influence and interactions may plausibly begin to get even more complex, resulting in a host of consequences that our current scientific machinery is ill-equipped to predict or understand.

In addition to their scale, speed, and complexity, IAs also challenge existing scientific processes and tools because of some of their potentially unique features that make research on IAs different—and arguably harder—than research on other socio-technical phenomena. IAs can be difficult to cleanly delineate compared to more discrete systems, and their scope, interaction, and even co-dependence can be fuzzy. IAs also evolve dynamically in response to technology and human behavior, leading to nonlinear interactions among IAs' up-code and down-code. Such recursivity also means it is not clear whether—to borrow from antiquity—we can ever step into the same IA twice. And all of this is confounded further by the opacity of many IAs; they lack the transparency researchers ideally look for (though transparency in socio-technical systems can be a double-edged sword; Nguyen 2022). We hope to better understand IAs' normative nature: like many sociotechnical systems, IAs necessarily encode (even if only implicitly) and promote specific value systems. That is, IAs shape information creation, dissemination, and use at levels ranging from underlying infrastructure to daily user experience. Integrating research across these levels to understand the bigger picture and to guide interventions poses non-trivial technical, methodological, and even ethical challenges.

Considering how IAs present both old and new research problems, understanding them can seem overwhelming. But if trends hold, we need to significantly increase the speed, diversity, and value of research on IAs to increase the chances that we better understand them and can shape their design to benefit democratic societies. This will require creativity and adaptability to navigate their complexity, dynamics, opacity, and entanglement with human behaviors and values. Developing a systematic understanding and tools to meaningfully improve IAs in ethical ways will require confronting these unique properties head-on.

To help structure future IA research in a way that we hope will be useful and intelligible, we propose a tripartite analytic framework adapted from Naugle et al. (2023): *Predicting, Understanding, and Prescribing (PUP)* interventions for complex sociotechnical systems like Information Architectures. We believe this framework can help categorize and advance research on different questions related to IAs while helping to organize results in a systematic way that advances our collective knowledge. Further, the framework also helps underscore the need for creative experimentation and novel tools that will help us make progress across all three of those areas.

Predicting the evolution and impacts of Information Architectures

New advances in computational power, data availability, analytic techniques, and theory promise to make prediction more tractable for IAs (Almaatouq et al. 2024; Smaldino 2023). With a focus on creating predictive models, there may be new research approaches that allow us to better anticipate things like the trajectories of, emergence of behaviors within, and social impacts stemming from specific IAs. With these approaches, we could better identify key parameters and variables that are the hallmark of persistent and adaptive IAs, and better forecast the outcomes of IA competitions. Without ignoring the difficulty of prediction in the face of organized complexity—a difficulty that should not be underestimated—being able to better assess futures as possible, plausible, or probable could better help us to anticipate how IA competitions lead to a range of potential outcomes: dominance of one IA over others,

subordination of an IA to others, changes in assortment or polarization in or across IAs, the emergence of hybrid or syncretic IAs, and the impacts of IAs on the kinds and quality of information shaping a population or populations.

Some key research questions we have identified related to prediction and IAs are:

- What new research approaches can be developed to predict the social impacts and emergent behaviors of specific IAs?
- How can we identify the key variables that determine the persistence and "fitness" of IAs?
- How can predictive models be utilized to forecast the outcomes of IA competitions?
- What are the practical or theoretical limits to prediction in the context of IAs?
- How can we categorize futures as possible, plausible, or probable for IA competitions and their diverse outcomes?
- What methods and approaches are needed to model IAs with sufficient complexity and fidelity?
- How can we create effective ways to test predictions and validate models in IA research?
- What strategies can be used to rapidly learn from incorrect predictions in IA research?

Tackling these kinds of research questions is likely to require significant advances in methods and approaches so researchers can not only model IAs and their interactions with sufficient complexity (Grimm et al. 2005), but create new ways to test predictions that help validate and refine those models. In particular, research here demands creative thinking about how to gather feedback on predictions that can rapidly enable us to learn where and when we are wrong. Recent advances in using LLMs to simulate social networks of discursively-coherent agents offer a promising direction (Horton 2023; Kozłowski et al. 2023; Park et al. 2023), though caution should still be exercised before drawing conclusions from AI models that are ultimately based on reproducing patterns in published corpora (Mitchell & Krakauer 2023).

Understanding Information Architectures

Understanding might seem to be an obvious priority for outlining a research agenda for IAs, but the organized complexity of IAs presents significantly different challenges than predicting their characteristics and dynamics. There is work needed for understanding their causal mechanisms, their characteristic features, and their informational boundaries. Prediction entails getting the right answers, understanding entails asking the right questions.

To better understand what IAs are and how they work, we need to tackle key research questions that involve theoretical and empirical research, including:

- How do we identify/draw boundaries around Information Architectures? What defines the scope and limits of an IA?
- What drives competition among IAs? What are the core fitness functions or units of selection for IAs?

- Are there inherent chaotic dynamics within IAs that could restrict our ability to understand and influence IAs?
- Are there fundamental limits to our ability to fully know an IA's structure and boundaries? Are there impossible IAs that we can exclude on principle?
- When does something qualify as an IA? Is the demarcation primarily structural or functional? Are there useful distinctions between “true” IAs and other related structures in terms of real-world impacts?
- What is the relationship between IAs and the actors that create/propagate them? To what extent are IAs constrained by their core media (smartphones, SMS, DMs, etc.)?

Understanding IAs will require a significantly funded and coordinated research campaign if we are to build a mature science in this space: a community of researchers, a common infrastructure with standardized terms and measures, a body of confirmatory research and validated theories, and an engagement with well-intentioned policy makers to apply this understanding.

Prescribing Interventions for Information Architectures

Discussing applications of understanding leads to the third need for IA research focused on prescription: that is, designing interventions to create mechanisms for coordination and establishing collectively desirable outcomes. IAs pose distinct challenges here because of their influence on key human characteristics like health, wealth, and security, not to mention emotion, morality, and social identity. This means that prescription research will also need to navigate these challenges within and across different societal norms and regulations, especially as those norms and regulations are likely to vary across political and cultural contexts. Interventions in complex social systems are fraught with peril due to the inevitability of unintended consequences, but lack of action may be unacceptable if the alternative outcomes are worse. While we anticipate some discomfort around the topic of prescribing interventions for IAs (and feel some of that discomfort ourselves), IAs are already being intervened upon, so understanding the consequences of those interventions seems prudent.

Some of the key questions we need to answer to effectively influence, engineer, design, and control IAs for collectively beneficial outcomes include:

- What creative coordination mechanisms can we use beyond top-down government regulation? How can we harness collective intelligence and market dynamics? What other design mechanisms can be explored?
- What are the rationales for proposed interventions succeeding in shaping IAs? What would we measure to determine impact or direction of effects of interventions?
- How can we establish collectively desirable IA outcomes? What are the critical tradeoffs among participants involved in these outcomes?
- How can we extrapolate from historical examples of IA transitions and interventions to prescribe future actions? How can we use comparative analysis of IA competition and evolutionary scenarios to identify key intervention points?

- How can we map, manage, and/or design incentives that reflect the complexity of IAs and their constituent parts (which will increasingly also include capable machines, as agents, that respond to rewards, policies, and utility functions)?

As with the other research foci, addressing these questions will require creative approaches that help navigate the practical, theoretical, and ethical challenges that each question presents individually as well as collectively. Tools that can improve counterfactual reasoning, which is necessary to pose and answer “what if?” questions, include simulations for predicting how IAs will work under different scenarios, as well as causal inference and mechanistic models. The irony is not lost on us that we will also almost certainly need the very computational capabilities that are now increasing the complexity of IAs in order to better predict, understand, and prescribe for them going forward.

8. Conclusions

As our social and information systems become more deeply entwined, the study of information architectures will be a tool for understanding and potentially shaping the techno-social environment that influences our lives. This paper makes an attempt to define, illustrate, and characterize the key dimensions of IAs.

By delving into the complexity, opacity, dynamism, and co-evolution of IAs with human behavior, we can avoid the perception of being a victim to technological change. Instead, we can gain insights into the underlying forces driving transformation and identify opportunities for human agency in shaping the development of and interaction between IAs. The scope of this research will require innovative approaches, interdisciplinary collaboration, and a recognition of its significant challenges. We covered vignettes and historical examples that offer some simple clues to how to understand and research IAs, and proposed a research framework focused on three key areas: prediction, understanding, and prescription. This framework was chosen to emphasize the proactive nature of the research agenda. By predicting the evolution and impacts of IAs, understanding their intricate workings and boundaries, and prescribing interventions for coordination and beneficial outcomes, we can perhaps steer society to adopt and evolve IAs that align with our collective values and aspirations. It is essential that we study not only individual online services and platforms, but also frame the larger problem of IAs as a critical research domain. The influence of IAs extends beyond any single application, shaping our economies, politics, cultures, and social interactions in profound ways.

Navigating ethical challenges and diverse societal norms involved in shaping IAs will require ongoing dialogue, experimentation, and adaptation. By working together across disciplines and sectors, we can strive to create IAs that promote human flourishing and collective well-being. The study of information architectures represents a critical step in this journey, empowering us to understand and shape the techno-social fabric that increasingly defines our world.

Acknowledgments

This paper arose from discussions at a workshop titled “Information Architectures as Sociotechnical Competitions,” held at the Santa Fe Institute in May, 2023 and organized by Paul Smaldino, Adam Russell, and Dan Patt. The workshop and subsequent research was funded by the Templeton World Charity (TWCF0715) and the Santa Fe Institute. We thank Amir Goldberg, Sandra González-Bailón, Haifeng Huang, and Alejandro Pérez Velilla for discussions that contributed to the ideas in this paper.

References

- Acemoglu, D., & Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6), 1488-1542.
- Acerbi, A. (2019). *Cultural evolution in the digital age*. Oxford University Press.
- Ahearn, L. M. (2001). *Invitations to love: Literacy, love letters, and social change in Nepal*. University of Michigan Press.
- Aker, J. C. 2010. Information from Markets Near and Far: Mobile Phones and Agricultural Markets in Niger. *American Economic Journal: Applied Economics* 2(3):46–59.
- Almaatouq, A., Griffiths, T. L., Suchow, J. W., Whiting, M. E., Evans, J., & Watts, D. J. (2024). Beyond playing 20 questions with nature: Integrative experiment design in the social and behavioral sciences. *Behavioral and Brain Sciences*, 47, e33.
- Anderson, P. W. (1972). More Is Different: Broken symmetry and the nature of the hierarchical structure of science. *Science*, 177(4047), 393-396.
- Antheunis, M. L., Tates, K., & Nieboer, T. E. (2013). Patients’ and health professionals’ use of social media in health care: motives, barriers and expectations. *Patient Education and Counseling*, 92(3), 426-431.
- Aral, S. (2020). *The hype machine: How social media disrupts our elections, our economy, and our health—and how we must adapt*. Currency.
- Arango, J. (2011). Architectures. *Journal of Information Architecture*, 3(1), 41-47.
- Arendt, H. (1951). *The origins of totalitarianism*. Harcourt, Inc.
- Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal*, 99(394), 116-131.
- Atkisson, C., Górski, P. J., Jackson, M. O., Hołyst, J. A., & D'Souza, R. M. (2020). Why understanding multiplex social network structuring processes will help us better understand the evolution of human behavior. *Evolutionary Anthropology*, 29(3), 102-107.
- Bakshy, E., Rosenn, I., Marlow, C., & Adamic, L. (2012). The role of social networks in information diffusion. In *Proceedings of the 21st International Conference on World Wide Web* (pp. 519-528).

- Boettke, P. J., Coyne, C. J., & Leeson, P. T. (2004). The many faces of the market. *Journal des Economistes et des Etudes Humaines*, 14(2), 5.
- Bozdog, E. (2013). Bias in algorithmic filtering and personalization. *Ethics and Information Technology*, 15, 209-227.
- Bradford, A. (2020) *The Brussels effect: How the European Union rules the world*.
- Bradford, A. (2023) *Digital Empires: The Global Battle to Regulate Technology*.
- Brady, W. J., Jackson, J. C., Lindström, B., & Crockett, M. J. (2023). Algorithm-mediated social learning in online social networks. *Trends in Cognitive Sciences*, 27(10), 947–960.
- Brancheau, J. C., & Wetherbe, J. C. (1986). Information architectures: methods and practice. *Information Processing & Management*, 22(6), 453-463.
- Brand, C. O., Mesoudi, A., & Smaldino, P. E. (2021). Analogy as a catalyst for cumulative cultural evolution. *Trends in Cognitive Sciences*, 25(6), 450-461.
- Brinkmann, L., Baumann, F., Bonnefon, J. F., Derex, M., Müller, T. F., Nussberger, A. M., ... & Rahwan, I. (2023). Machine culture. *Nature Human Behaviour*, 7(11), 1855-1868.
- Brunton, F. (2020). *Digital cash: The unknown history of the anarchists, utopians, and technologists who created cryptocurrency*. Princeton University Press.
- Centola, D. (2015). The social origins of networks and diffusion. *American Journal of Sociology*, 120(5), 1295-1338.
- Centola, D., & Baronchelli, A. (2015). The spontaneous emergence of conventions: An experimental study of cultural evolution. *Proceedings of the National Academy of Sciences*, 112, 1989-1994,
- Chan, H. F., Önder, A. S., Schweitzer, S., & Torgler, B. (2023). Twitter and citations. *Economics Letters*, 231, 111270.
- Coward, L. A. (2014). Brain computational primitives. *Procedia Computer Science*, 41, 164–175.
- Culkin, J. M. (1967, March 18). A Schoolman's Guide to Marshall McLuhan. *The Saturday Review*, 51–53, 66–79.
- David, P. A. (1985). Clio and the Economics of QWERTY. *The American Economic Review*, 75(2), 332-337.
- Dawkins, R. (1986). *The blind watchmaker: Why the evidence of evolution reveals a universe without design*. WW Norton & Company.
- De Vos, P. S. (2006). Research, development, and empire: State support of science in the later Spanish Empire. *Colonial Latin American Review*, 15(01), 55-79.
- Dillon, A., & Turnbull, D. (2005). Information architecture. *Encyclopedia of library and information science*, 1-9.
- DiResta, R. (2024) *Invisible rulers: The people who turn lies into reality*. PublicAffairs.
- Doctorow, C. (2020) *How to Destroy Surveillance Capitalism*. Medium Editions.

- Doctorow, C. (2023, January 23). The 'enshittification' of TikTok. *Wired*.
<https://www.wired.com/story/tiktok-platforms-cory-doctorow/>
- Donath, J. (2014). *The Social Machine: Designs for Living Online*. Cambridge, MA, MIT Press.
- Downey, L., & Banerjee, S. (2011). Building an information architecture checklist. *Journal of Information Architecture*, 2(2).
- Drelichman, M., Vidal-Robert, J., & Voth, H. J. (2021). The long-run effects of religious persecution: Evidence from the Spanish Inquisition. *Proceedings of the National Academy of Sciences*, 118(33), e2022881118.
- Epstein, B. (2015). *The ant trap: Rebuilding the foundations of the social sciences*. Oxford University Press.
- Farrell, H., & Newman, A. L. (2019). Weaponized interdependence: How global economic networks shape state coercion. *International Security*, 44(1), 42-79.
- Farrell, H., & Newman, A. (2023). *Underground empire: How America weaponized the world economy*. Henry Holt and Company.
- Feezell, J. T., Wagner, J. K., & Conroy, M. (2021). Exploring the effects of algorithm-driven news sources on political behavior and polarization. *Computers in Human Behavior*, 116, 106626.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Psychology Press.
- Gillespie, T. (2018). *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media*. Yale University Press.
- González-Bailón, S., Lazer, D., Barberá, P., Zhang, M., Allcott, H., Brown, T., ... & Tucker, J. A. (2023). Asymmetric ideological segregation in exposure to political news on Facebook. *Science*, 381, 392-398.
- Gorwa, R. (2024). *The Politics of Platform Regulation: How Governments Shape Online Content Moderation*. Oxford University Press.
- Granovetter, M. S. (1973). The strength of weak ties. *American Journal of Sociology*, 78(6), 1360-1380.
- Grimm, V., Revilla, E., Berger, U., Jeltsch, F., Mooij, W. M., Railsback, S. F., ... & DeAngelis, D. L. (2005). Pattern-oriented modeling of agent-based complex systems: lessons from ecology. *Science*, 310(5750), 987-991.
- Guilbeault, D., Baronchelli, A., & Centola, D. (2021). Experimental evidence for scale-induced category convergence across populations. *Nature Communications*, 12, 327.
- Hanson, R. (2007). Logarithmic markets scoring rules for modular combinatorial information aggregation. *The Journal of Prediction Markets*, 1(1), 3-15.
- Headrick DR, Griset P. (2001) Submarine Telegraph Cables: Business and Politics, 1838–1939. *Business History Review*. 75(3):543-578.

- Henrich, J. (2004). Cultural group selection, coevolutionary processes and large-scale cooperation. *Journal of Economic Behavior & Organization*, 53(1), 3-35.
- Henrich, J. (2015). *The secret of our success: How culture is driving human evolution, domesticating our species, and making us smarter*. Princeton University press.
- Henrich, J., McElreath, R. (2007) Dual inheritance theory: The evolution of human cultural capacities and cultural evolution. In *Oxford Handbook of Evolutionary Psychology*, edited by Robin Dunbar and Louise Barrett, 555-570. Oxford: Oxford University Press
- Herman, E. S., & Chomsky, N. (1988). *Manufacturing consent: The political economy of the mass media*. Pantheon.
- Hess, C. and E. Ostrom. (2003) Ideas, Artifacts, and Facilities: Information as a Common-Pool Resource. *Law and Contemporary Problems*, V 66, No. 1/2, pp. 111-145.
- Hess, C. and E. Ostrom. (2007) Introduction: An Overview of the Knowledge Commons. In *Understanding Knowledge as a Commons: From Theory to Practice*. C. Hess and E. Ostrom, eds.
- Heyd, T., & Puschmann, C. (2017). Hashtagging and functional shift: Adaptation and appropriation of the #. *Journal of Pragmatics*, 116, 51-63.
- Hickey, D., Schmitz, M., Fessler, D., Smaldino, P. E., Muric, G., & Burghardt, K. (2023, June). Auditing Elon Musk's impact on hate speech and bots. In *Proceedings of the international AAAI conference on web and social media* (Vol. 17, pp. 1133-1137).
- Horton, J. J. (2023). Large language models as simulated economic agents: What can we learn from *homo silicus*? (No. w31122). *National Bureau of Economic Research*.
- Hruschka, D. J. (2010). *Friendship: Development, ecology, and evolution of a relationship*. University of California Press.
- Ingram, D. (2024, March 22) Fewer people are using Elon Musk's X as the platform struggles to attract and keep users, according to analysts. *NBC News*.
<https://www.nbcnews.com/tech/tech-news/fewer-people-using-elon-musks-x-struggles-keep-users-rcna144115>
- Jara-Figueroa, C., Yu, A. Z., & Hidalgo, C. A. (2019). How the medium shapes the message: Printing and the rise of the arts and sciences. *PLOS ONE*, 14(2), e0205771.
- Jensen, R. (2007) The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector, *The Quarterly Journal of Economics*, 122(3), 879–924, <https://doi.org/10.1162/qjec.122.3.879>
- Jones, M. L. (2020). Cookies: a legacy of controversy. *Internet Histories*, 4(1), 87-104.
- Kauffman, S. A. (2000). *Investigations*. Oxford University Press.
- Kim, M., & Cha, J. (2017). A comparison of Facebook, Twitter, and LinkedIn: Examining motivations and network externalities for the use of social networking sites. *First Monday*.
- Kozlowski, A. C., Kwon, H., & Evans, J. A. (2024). In Silico Sociology: Forecasting COVID-19 Polarization with Large Language Models. *arXiv preprint arXiv:2407.11190*.

- Krohn, R., & Weninger, T. (2022, May). Subreddit links drive community creation and user engagement on Reddit. In *Proceedings of the International AAAI Conference on Web and Social Media* (Vol. 16, pp. 536-547).
- Lampel, J., & Bhalla, A. (2007) The Role of Status Seeking in Online Communities: Giving the Gift of Experience, *Journal of Computer-Mediated Communication*, 12(2), 434–455.
- Ledwich, M., Zaitsev, A., & Laukemper, A. (2022). Radical bubbles on YouTube? Revisiting algorithmic extremism with personalised recommendations. *First Monday*.
- LeFebvre, L. E., Allen, M., Rasner, R. D., Garstad, S., Wilms, A., & Parrish, C. (2019). Ghosting in emerging adults' romantic relationships: The digital dissolution disappearance strategy. *Imagination, Cognition and Personality*, 39(2), 125-150.
- Lelkes, Y., Sood, G., & Iyengar, S. (2017). The hostile audience: The effect of access to broadband internet on partisan affect. *American Journal of Political Science*, 61(1), 5-20.
- Lerman, K., Feldman, D., He, Z., & Rao, A. (2024). Affective polarization and dynamics of information spread in online networks. *npj Complexity*, 1(1), 8.
- Luccioni, A. S., Viguier, S., & Ligozat, A. L. (2023). Estimating the carbon footprint of bloom, a 176b parameter language model. *Journal of Machine Learning Research*, 24(253), 1-15.
- March, J. G., & Simon, H. A. (1993). *Organizations*. Wiley.
- McClain, C., Anderson, M., & Gelles-Watnick (2024, June 12) How Americans Navigate Politics on TikTok, X, Facebook and Instagram. *Pew Research Center*.
<https://www.pewresearch.org/2024/06/12/how-americans-navigate-politics-on-tiktok-x-facebook-and-instagram/>
- McLuhan, M. (1964) *Understanding Media: The Extensions of Man*. McGraw-Hill.
- Merton, R. K. (1970) *Science, Technology and Society in 17th Century England* (New Jersey: Humanities Press)
- Milzman, J., & Moser, C. (2023). Decentralized core-periphery structure in social networks accelerates cultural innovation in agent-based modeling. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems* (pp. 2213-2221).
- Mitchell, M., & Krakauer, D. C. (2023). The debate over understanding in AI's large language models. *Proceedings of the National Academy of Sciences*, 120(13), e2215907120.
- Morville, P., & Rosenfeld, L. (2006) *Information architecture for the world wide web*. O'Reilly.
- Munn, L. (2020). Angry by design: Toxic communication and technical architectures. *Humanities and Social Sciences Communications*, 7(1), 52.
- Naugle, A., Russell, A., Lakkaraju, K., Swiler, L., Verzi, S., & Romero, V. (2023). The Ground Truth program: simulations as test beds for social science research methods. *Computational and Mathematical Organization Theory*, 29(1), 1-19.
- Newson, L., & Richerson, P. J. (2021). *A story of us: a new look at human evolution*. Oxford University Press.

- Nguyen, C. T. (2022). Transparency is surveillance. *Philosophy and Phenomenological Research*, 105(2), 331-361.
- Nikolov, D., Lalmas, M., Flammini, A., & Menczer, F. (2019). Quantifying biases in online information exposure. *Journal of the Association for Information Science and Technology*, 70(3), 218-229.
- Norman, D. (1990). *The design of everyday things*. Basic books.
- Park, J. S., O'Brien, J., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (pp. 1-22).
- Paul, C., & Matthews, M. (2016) *The Russian "firehose of falsehood" propaganda model, why it might work and options to counter it*. RAND Corporation.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. Morgan Kaufmann.
- Quito, A. (2019, April 16). Jack Dorsey admits Twitter's "like" button should never have existed. *Quartz at Work*. <https://qz.com/work/1597035/ted2019-jack-dorsey-rethinks-twitthers-like-feature>
- Ribeiro, M. H., Ottoni, R., West, R., Almeida, V. A., & Meira Jr, W. (2020, January). Auditing radicalization pathways on YouTube. In *Proceedings of the 2020 conference on fairness, accountability, and transparency* (pp. 131-141).
- Richerson, P., Baldini, R., Bell, A. V., Demps, K., Frost, K., Hillis, V., ... & Zefferman, M. (2016). Cultural group selection plays an essential role in explaining human cooperation: A sketch of the evidence. *Behavioral and Brain Sciences*, 39, e30.
- Rosenfeld, M. J., Thomas, R. J., & Hausen, S. (2019). Disintermediating your friends: How online dating in the United States displaces other ways of meeting. *Proceedings of the National Academy of Sciences*, 116(36), 17753-17758.
- Ruegg, W. (1996) *A History of the University in Europe* (Vol. 2). Cambridge University Press,.
- Russell, C. J. S., & Muthukrishna, M. (2021). Dual inheritance theory. In T.K. Shackelford, V. A. Weekes-Shackelford (eds.), *Encyclopedia of Evolutionary Psychological Science* (pp. 2140-2146). Springer Nature.
- Sáenz-Royo, C., Gracia-Lázaro, C., & Moreno, Y. (2015). The role of the organization structure in the diffusion of innovations. *PLOS ONE*, 10(5), e0126076.
- Scott, J. C. (1998). *Seeing like a state: How certain schemes to improve the human condition have failed*. Yale University Press.
- Scott, D. A., Valley, B., & Simecka, B. A. (2017). Mental health concerns in the digital age. *International Journal of Mental Health and Addiction*, 15, 604-613.
- Sethi, R., Seager, J., Cai, E., Benjamin, D. M., & Morstatter, F. (2021). Models, Markets, and the Forecasting of Elections. *arXiv preprint arXiv:2102.04936*.

- Shapiro JN, Weidmann NB. Is the Phone Mightier Than the Sword? Cellphones and Insurgent Violence in Iraq. *International Organization*. 2015;69(2):247-274.
doi:10.1017/S0020818314000423
- Shapiro, B. J. (1971). The universities and science in seventeenth century England. *Journal of British Studies*, 10(2), 47-82.
- Shapiro, S. (2023) *Fancy Bear Goes Phishing: The Dark History of the Information Age in Five Extraordinary Hacks*.
- Smaldino, P. E. (2014). The cultural evolution of emergent group-level traits. *Behavioral and Brain Sciences*, 37(3), 243-254.
- Smaldino, P. E. (2023). *Modeling social behavior: mathematical and agent-based models of social dynamics and cultural evolution*. Princeton University Press.
- Smaldino, P. E., & Jones, J. H. (2021). Coupled dynamics of behaviour and disease contagion among antagonistic groups. *Evolutionary Human Sciences*, 3, e28.
- Smaldino, P. E., & Richerson, P. J. (2013). Human cumulative cultural evolution as a form of distributed computation. In *Handbook of Human Computation* (pp. 979-992). New York, NY: Springer New York.
- Sperber, D. (1994). The modularity of thought and the epidemiology of representations. *Mapping the mind: Domain specificity in cognition and culture*, 39-67.
- Spiekermann, S., A. Acquisti, et al. (2015). The challenges of personal data markets and privacy. *Electronic markets*, 25, 161-167.
- Srba, I., Moro, R., Tomlein, M., Pecher, B., Simko, J., Stefancova, E., ... & Bielikova, M. (2023). Auditing YouTube's recommendation algorithm for misinformation filter bubbles. *ACM Transactions on Recommender Systems*, 1(1), 1-33.
- Steiglechner, P., Smaldino, P. E., Moser, D., & Merico, A. (2023). Social identity bias and communication network clustering interact to shape patterns of opinion dynamics. *Journal of the Royal Society Interface*, 20(209), 20230372.
- Steinwender, C. (2018). Real effects of information frictions: When the states and the kingdom became united. *American Economic Review*, 108 (3): 657-96.
- Sumter, S. R., & Vandenbosch, L. (2019). Dating gone mobile: Demographic and personality-based correlates of using smartphone-based dating applications among emerging adults. *New Media & Society*, 21(3), 655-673.
- Thicke, M. (2018). Market epistemology. *Synthese*, 195(12), 5571-5594.
- Timmermans, E., Hermans, A.-M., & Oprea, S. J. (2021). Gone with the wind: Exploring mobile daters' ghosting experiences. *Journal of Social and Personal Relationships*, 38(2), 783-801.
- Turner, M. A., Singleton, A. L., Harris, M. J., Harryman, I., Lopez, C. A., Arthur, R. F., ... & Jones, J. H. (2023). Minority-group incubators and majority-group reservoirs support the diffusion of climate change adaptations. *Philosophical Transactions of the Royal Society B*, 378(1889), 20220401.

- Viégas, F. B., Golder, S., & Donath, J. (2006, April). Visualizing email content: portraying relationships from conversational histories. In *Proceedings of the SIGCHI conference on Human Factors in computing systems* (pp. 979-988).
- Wagner, C., Strohmaier, M., Olteanu, A., Kıcıman, E., Contractor, N., & Eliassi-Rad, T. (2021). Measuring algorithmically infused societies. *Nature*, 595, 197-204.
- Watts, D. J. (1999). Networks, dynamics, and the small-world phenomenon. *American Journal of Sociology*, 105(2), 493-527.
- Weaver, W. (1948). Science and complexity. *American Scientist*, 36(4), 536-544.
- Wiener, N. (1950). *The human use of human beings: Cybernetics and society*. Houghton Mifflin.
- Wu, T. 2016. *The Attention Merchants: The Epic Struggle to Get Inside Our Heads*.
- Young, H. P. (2006). The diffusion of innovations in social networks. *The economy as an evolving complex system III: Current perspectives and future directions*, 267, 39.
- Zegart, A. B. (2022). *Spies, Lies, and Algorithms: The History and Future of American Intelligence*. Princeton University Press.
- Zollman, K. J. (2013). Network epistemology: Communication in epistemic communities. *Philosophy Compass*, 8(1), 15-27.
- Zribi, W., Boufateh, T., & Guesmi, K. (2023). Climate uncertainty effects on bitcoin ecological footprint through cryptocurrency environmental attention. *Finance Research Letters*, 58, 104584.