Interop

The Promise and Perils of
Highly Interconnected Systems

John Palfrey
and
Urs Gasser

BASIC BOOKS
A MEMBER OF THE PERSEUS BOOKS GROUP
New York

INTRODUCTION

Why Interop Matters

The Internet has made the world at once a smaller and a more complex place. Digital technologies connect billions of people, businesses, organizations, and governments with each other in ways that enhance our lives but that we don't fully understand. We are interconnected as never before, to our enormous benefit: we stay in touch with far-away family and friends for low cost; we learn about news instantaneously, access knowledge remotely, collaborate more efficiently, and do all kinds of business online. Our most complex systems—government agencies, financial institutions, transportation infrastructures, health care and energy systems—are linked by these new, invisible information channels, which are essential components of today's global economy.

This capacity for connection is about more than just making our lives more convenient or efficient. Organizations can become more specialized, better at what they do, and more able to collaborate effectively across distance and time—whether in business, public life, or civic activism—in ways that are changing the course of history. Consider what a highly connected network of people in the Middle East, many of them very young,

were able to accomplish, with little in the way of central coordination, during the Arab spring of 2011. They toppled long-standing regimes, one after another, through peaceful activism that was powered by their high degree of digital connectedness.

But this growing level of interconnectedness comes at an increasingly high price. We make big trade-offs as we become digitally connected everywhere and anytime. We struggle to keep up with overflowing e-mail inboxes; we feel overwhelmed by the flood of news and information coming at us from all directions. We are also more vulnerable, in ways that are less obvious and less well understood. The same infrastructure that enables us to create, store, and share information can put our privacy and security at risk. Data breaches and privacy invasions make the news daily, illustrating what can happen when massive amounts of data are exchanged among complex systems without adequate safeguards. These risks are omnipresent in the digital age. They touch every aspect of modern life through the information exchanged with banks, credit card companies, mobile phone carriers, tax authorities, entertainment giants, or online businesses.

The problems of too much interconnectivity present enormous challenges both for organizations and for society at large. Our most advanced systems and infrastructures have become so complex that they are hard to manage effectively. Our financial system, for instance, has fallen into deep crisis due in part to the new vulnerabilities, complexities, and domino effects resulting from unprecedented digital connectivity. Our technological networks are so pervasive, and we use them so intensely, that we have good reason to worry that data about ourselves and our families might float out of our sight and our control. In such cases, the problem is not the interconnectivity itself but, rather, the fact that it is not adequately checked or managed.

In other cases, we suffer from too little connectivity. We struggle, for instance, to reform our health care system in no small part because we cannot get our information systems to work together properly with one another. Higher degrees of connectivity and information sharing among our health care providers would make the health care system vastly more efficient and

effective at providing care—and yet that connectivity eludes us. If the devices we use and the households we live in could "talk" to each other, we could dramatically reduce our energy consumption by creating a "smart grid" that would allow for efficient energy supply. More often than not, our future success in addressing the big societal challenges of our time, from health care to climate change, will depend heavily on our ability to create better interfaces and connections among complex systems and our ability to share information appropriately.

This challenge—creating better, more useful connectivity while simultaneously finding better ways to manage its inherent risks—inspired this book. As societies, we have rushed to build information and communications infrastructures that enhance connectivity and enable the flow of information among individuals, organizations, and systems. But we have not yet developed a normative theory identifying what we want out of all this interconnectivity. We call this theory *interoperability*, or *interop*. The payoff from our theory of interop is that it can help us decide where we need interconnectedness in complex systems and at what level—and where we don't. Without such a theory, we lack a stable framework for figuring out how to harness the benefits of the digital technologies that connect us while still protecting our core societal values. And we have not yet refined a sense of which tools will get us to optimal levels of interop. This book has been written to help meet these urgent challenges—challenges that are at once highly conceptual and deeply practical.

The main purpose of the theory of interoperability is to help define the optimal level of interconnectedness and to lay out a path for achieving it. As a first step, we must develop a new lens for analyzing how complex systems, components, and applications are connected—or sometimes, inexplicably, still separated. Second, we must take a deep look at the promises and the drawbacks that come with increased connectivity. We need to balance the costs and the benefits of the connectivity we create, both in the short and long terms. We can then assess how much interconnectedness we should aim to achieve among our institutions, systems, and peoples. Most important, a theory of interoperability leads to a clear understanding

of the mechanisms—technical, organizational, or legal—through which interop can be achieved and shows how we might optimize the interoperability levels of complex systems.

This book explores in depth two interop problems. The first is to figure out how to define and get to an optimal level of interoperability in complex systems. The second is to deal with the adverse effects of interoperability: loss of diversity, increasingly pressing concerns about its effects on individual privacy and security, and the risk of locking in older technologies and hindering innovation. An enormous amount hangs in the balance. Our economies, our personal well-being, and our environment will all be affected by whether these two interop problems can be solved in our most critical complex systems.

We, the authors, have been studying interoperability through a variety of methods for the past decade. We are both law professors and researchers, interested in the way the ongoing explosion of information technologies affects societies around the world. Our joint research project started out as a transatlantic collaboration. One team was based in Europe, at a leading research center at the University of St. Gallen in Switzerland. The other was based at the Berkman Center for Internet and Society at Harvard University. Our methodology is based on the development and analysis of a broad range of case studies. We started out with a series of cases that relate primarily to information and communications technologies, which is our core field of study. We have talked to hundreds of people and hosted workshops on three continents with experts in a wide range of fields.

As we got deeper and deeper into the topic, we began to see the reach of interop beyond the context of our core field. We began to research fields outside our own: economics, business, systems theory, psychology. Through our case studies, drafted by members of an interdisciplinary team of researchers at our respective centers, we began to examine areas farther afield where interop matters greatly.

Some of the biggest challenges of the age are in fact interop problems. Consider health care reform, which relies upon getting interoperability

right in the context of electronic health records, or climate change, which turns in part on the emergence of a next-generation energy delivery system, the smart grid. We have posted both of these case studies online, free for anyone to read, as a companion to this book.[1] We present them as the raw data from which we have built the theory of interop and the practical suggestions that we offer in this book. Our idea, in the spirit of transparency, is that anyone can look at the data from which we have drawn our conclusions, and we hope to provoke dialogue on these pressing issues.

Much of our research has involved conducting in-depth interviews and convening workshops with experts in the fields of computing, law, and psychology, as well as in many fields of industry. Over the many years of conducting interviews on this topic, we have never found a single person who thinks that interop is anything other than a good thing in general. That is the starting point: people generally want higher levels of interconnection. After that, there is not a lot of common ground. There is no single, agreed-upon definition of interoperability. There are many views about what interop is and how it should be achieved. And there are even more views about how, if at all, the problems to which interop gives rise should be addressed.

There is no one-size-fits-all definition of interoperability. In the most general sense, in the context of information technologies, interoperability is the ability to transfer and render useful data and other information across systems, applications, or components. But it is important to go beyond this core understanding to explore a broader understanding of what interop means in different contexts and at different levels.

In theoretical terms, interoperability functions on four broad layers of complex systems. Understanding this structure is essential to understanding how interop works and how society ought to go about achieving (or thwarting) it. Interop is not just about the flow of data or about technology; it involves essential questions of human and institutional interaction as well. The problems associated with interop are just as much about culture as they are about technology.

The first layer is technological. Think of the hardware and the code in computing systems or the train tracks in the transportation systems. Interoperability at this layer means that, in the most basic sense, the systems can connect to one another, often through an explicit, agreed-upon interface.

The second layer of interoperability is the data layer. This layer is closely paired with the technology layer; indeed, the two are often inextricably linked. It is not enough for the technological systems to be able to exchange signals or to pass material from one to the next. If the receiving party cannot understand the data, then the technological interoperability is worthless. Imagine that you receive an e-mail with an attachment on your smartphone. When you click on the attachment, you get an error message: you can't read that attachment on your device. In this case, the software on your smartphone can receive the message but cannot render the data useful to you.

The third layer of interoperability is the human layer. This layer is much more abstract than the technology and data layers, but it is very important to the success of interoperability. It is one thing for the e-mail systems to exchange messages between them and for the data to be passed successfully across those interoperable systems, but it is another thing for the humans at either side of the exchange of information to understand each other and to act upon that exchange. Language is one way to think about the human layer of interoperability—in order to communicate, we need a common language—but that is only the starting point. We also need to be able to work together in other ways. Interop often succeeds or fails based on whether we are willing to put effort into working together as human beings.

At the highest and most abstract layer, we consider interoperability at the institutional layer. Just as it is essential that people work together, it is also frequently important that societal systems engage effectively. The legal system is one example of an institutional layer of interoperability (or its absence). For instance, if two companies in different countries want to collaborate, they must share a common understanding of, say, contract law. Likewise, if two companies want to start a joint venture, they need a shared understanding of corporate law. This does not mean that the two countries need to have identical legal systems or that the two companies need to

adopt the same internal processes or rules. They only need to have *enough* in common that the interests of both are protected. Interoperability at the institutional layer makes possible high levels of collaboration and exchange without making the parties identical.

Given the importance of each of these four layers, no short definition of interop ends up being particularly satisfying. Interoperability is highly context specific. And so, rather than aiming for a single definition that can apply across different sectors and cases, we consider the specific contours of the interoperability at work in each example, across the four layers; we operate pragmatically and with an open working definition. This process approach to defining interoperability is meant to avoid prejudging the best way to accomplish interoperability. It is also intended to reflect the idea that interop is not a binary concept. There are degrees and types of interop, which fall along a multidimensional spectrum.

Interop also means different things to different people. The kind of interop that matters to computer users—whether an e-mail comes through the system legibly, for example—may be different from the kind of interop that matters to the Internet service providers who have to send the messages, to the companies that make the software and hardware that make the systems work in the first place, and to the police who from time to time want to be able to intercept those e-mails in order to apprehend a criminal. In the context of signing up for a new social network online, interoperability might mean being able to sign into one program or website (such as Twitter or Facebook) and having personal information seamlessly and securely transferred as needed to a variety of merchants and service providers (such as the mobile apps foursquare or SCVNGR). In the context of online music, recording industry executives might view interoperability as being able to sell their content securely through a variety of online channels and have it play on many approved devices. Web service and mashup platform providers care about seamless data transmission and easy extension and integration of data sources by users and small developers.[2]

The point is that different people and firms will have different perspectives on what interoperability means in a given context, how much interoperability is optimal, and how it ought to be accomplished. The incentives

related to interoperability can vary greatly. Some firms will seek to use interoperability to keep people within their systems; others will want to profit directly or indirectly by enabling others to innovate on the basis of an open, broadly interoperable platform.

Interoperability should be an explicit goal in national and international discussions of business, law, and policy because the upsides of interoperability are massive: it fosters innovation and competition, enhances diversity, gives consumers choice, and can lead to unexpected benefits over time. Interoperability is not an end in itself; rather, it is a means to accomplish other societal goals, such as growing the economy, fighting climate change, and improving the quality of health care. Our goal should be to harness the great potential of interoperability while avoiding some of its possible downsides.

Interop can help many people in many contexts. For instance, consumers who want to be able to choose from a broad range of applications for their home entertainment systems are well advised to purchase a system that offers interoperability across different providers and services. Entrepreneurs who seek to develop and market their own web application are usually more likely to succeed if they pursue an interop-based approach. Business executives should usually strive for interoperability among teams, work flows, and the like within their organizations. Government agencies operate at lower costs and with greater efficiency, and thus can provide better service to citizens, when they (and their systems) work together. When legislators and policy makers are creating or adjusting legal frameworks aimed at fostering innovation, they should consider the various approaches described in this book to create incentives for increasing both technical and institutional interoperability.

Our approach to interoperability takes several forms. Although the most obvious context for the argument about the benefits of interop is the information and communications technology sector, interoperability matters in sectors throughout the economy. That said, although we have studied historical examples—transportation and finance in particular—to glean

insights into how society has built successful interoperable systems in the past, we focus here primarily on debates that rage today in the information and communications technologies sector. These debates relate to the distribution of digital music and movies, document formats, and the long-term preservation of human knowledge. We scan the horizon for issues that are just emerging, such as cloud computing,[3] the smart grid, e-health records, and online identity systems. And we make a series of arguments about how interoperability might be achieved through law, policy, technology, and innovations in the marketplace. As we look ahead, we contend that interop-related challenges will only grow harder to manage as our systems grow more complex and interconnected.

Historical perspective is an important starting point for our study of interop. Systems have failed to work together since time immemorial. And when they have succeeded, humans have sometimes had to live with unforeseen and unwanted consequences.

The American rail system is one example of how people have worked together to solve interop problems. In May 1869, in the middle of the United States, a small group witnessed the ceremony of the golden spike, a major event in the history of interoperability. The witnesses celebrated the connection of the railroad systems, which now reached from the Pacific Ocean to the Atlantic. The golden spike, driven into the earth at the center of the country, made it possible for a train to connect the two great oceans of the world. Civilizations could be joined in a new way from one part of the globe to another.[4]

On its simplest level, the connection of train tracks from one ocean to another, across the massive North American continent, is a story about a technology. The technology of train tracks and engines and cars is an essential infrastructure in a modern economy. This technology was developed not by one single government or one firm. It was paid for and built by a whole lot of people with many different financial, political, and social interests. Those many interests were not necessarily aligned. But somehow, a system emerged that made it possible to travel at high speed from the

edge of one ocean to another. The "somehow" was a commitment to make a system that could interoperate. The idea was not to develop a single train system that was the same in every respect everywhere or that was owned by the same people; rather, the goal was to achieve one that would work together across different owners and different plans for usage.

Interoperability of the rail system in the late nineteenth century in America made many other good things possible. This technology carried goods harvested or manufactured in one corner of the country to others, making new markets accessible. The railroad was also an early communications network. Over its tracks rode people, ideas, and cultural norms. The interconnected, interoperable railroad system made possible a newer, faster way for people to communicate, for markets to grow more complex and profitable, and for cultures to become connected to one another. The interoperability that the US rail system made possible reached far beyond the ability to convey trains from one place to another.

The twentieth century is full of examples that illustrate the importance of interoperability as a driver of innovation, growth, and benefits to consumers. The further development of the transportation infrastructure is one such example. Consider the ease with which people can travel by air or car across the countries of Europe, for instance, and the number of systems that need to work together to make such seamless—and safe—travel possible. Financial systems are an equally instructive example: the extent to which currency can flow from one jurisdiction to another has driven international trade and cross-cultural exchange of many sorts.

In neither of these large-scale examples—transportation and finance—has interop put an end to diversity. Systems must have sufficient overlap to work together, but they do not need to be completely standardized. This key distinction—between sameness and interoperability—recurs throughout the examples we explain in this book. In the best cases, even while systems and people are enabled to work together, the powerful force of diversity can be preserved. The point was not that there needed to be a single train company or a single bank that everyone had to use. Nor did all the train companies or banks have to do everything the same way. They just

had to agree to do *some* things in ways that would interoperate. Crucially, the things they agreed to do in an interoperable manner had to be the *right* things.

Think of the trains themselves. The sizes and shapes of the trains could vary significantly from one company to another. The ways they hired and staffed their train systems could be quite different. The policies for what a given company would carry across the tracks, and how, could be widely diverse. But the gauge of the tracks and the distance between the wheels of the trains, along with a few other technical specifications, had to be the same.

Currency is another example of the compatibility between diversity and interoperability. The Swiss franc coexists with the euro and the British pound in the regional economy of Europe. Each of these currencies coexists with the US dollar, the Chinese yuan, and dozens of other widely used currencies. And yet a global economy has emerged whereby people from each of these jurisdictions can trade together without major hassles. Enabling this trade was not a process of standardizing on a single currency, with a single value and a single governor. The process has been more subtle than that, emerging from the bottom up over a long period of time and in turn enabling local diversity while giving rise to a global system of finance. The system has been made to interoperate through the establishment of intermediaries, rules, and laws.

One of the primary benefits of interoperability is that it can preserve key elements of *diversity* while ensuring that systems work together in the ways that matter most. One of the tricks to the creation of interoperable systems is to determine what the optimal level of interoperability is: in what ways should the systems work together, and in what ways should they not?

The benefits of interoperability are vast. In particular, interoperable systems make all our lives easier. Interoperable systems can make us more efficient by lowering the costs of switching between and among varying tasks. They can afford consumers more choice by limiting the effects of being locked in to any one system. They can promote cross-cultural

understanding, the free movement of ideas, and the flow of trade. They can support a competitive environment for businesses. And they can often lead to innovation in the marketplace.

And yet interop is not always an unalloyed good. The breakdown in the global financial system in 2008 and the subsequent crisis stemming from Greece's defaults in 2010 illustrate the dangers of highly interoperable systems. The economies of Europe, in particular, are tightly interconnected: the European Union is an economic unit by design. The downside of this degree of interconnection is that Greece's debt woes have meant that Germany, France, and other economically stronger countries have had to foot the bill for much of Greece's overspending. The European Union, in turn, has become deeply linked to the economies of the United States and many big Asian markets. As a result, the effects of Greece's ongoing problems have been felt in every economy in the world. We have become very good at connecting our economies, but not especially good at isolating the problems that arise in one part of the world from the rest.

For a much simpler, prosaic example of a situation in which you want significant but not complete interoperability, consider a car trip. You are driving home to New York from a visit to Boston. You're tired and bleary-eyed. You also realize that you're short on gas: the little red gas indicator has lit up beside the odometer. You decide it would be a good idea to pull off the highway before driving much further.

At the next gas station you see, you put your credit card into the computer attached to the gas pump and lift the nozzle. You try to put the nozzle into your car's gas tank, but it won't go. You try again. It still won't go. You realize, to your surprise, that the nozzle is the wrong size. After a few more tries, and a curse or two, it dawns on you that you've made a mistake. Your car takes ordinary fuel, but the nozzle in your hand is connected to the tank of diesel fuel. You pick up another nozzle, corresponding to the correct kind of fuel, and soon you're on your way with a full tank.

Several things had to work together to ensure that you got your gas without ruining your engine. A complex financial system enabled you to enter your credit card information into a computing system associated with the pump at the gas station. After a flurry of bits made their way to your bank

and back, the company selling the gas decided it was safe to let the gas flow into your car, because the funds would be released by the credit card issuer. And years earlier, the manufacturer of your car sized the receptacle for gas correctly so as to allow the nozzle dispensing ordinary unleaded fuel to fit properly. The gas station had to offer nozzles that fit the corresponding apertures, and so forth. Interop is the secret to these complex systems working together to enable you to get gas efficiently and safely into your car.

At the same time, the system was set up so that you couldn't put the wrong fuel—in this case, diesel, which would have harmed your engine—in your car. Likewise, if you had introduced a stolen credit card, the financial system would have denied you the fuel. The system was designed to not interoperate when it wasn't meant to. In this case, the system was meant to correct for human error (your bleary-eyed reach for the diesel nozzle) and to prevent cheating (someone's attempt to spend another person's money on fuel).

We do not always want things to interoperate completely. Sometimes we want brakes on interoperability to correct against human error, as the diesel gas pump example demonstrates (purposeful noninteroperability). Other times, we want brakes to prevent fraud, as with the example of the credit card: interoperability is blocked if there is a possibility the card is stolen. We want to make sure that the parts of the system can always work together but also that the system can throw up roadblocks or speed bumps where necessary (limited, or conditional, interoperability).

The same principle holds true with all sorts of other complex systems. In the environment of the web, we want the system to be able to pass data from one place to another, but we also want it to be able to include brakes that stop the wrong kind of personal information (for instance, health-related data) from flowing from one place to another in the wrong cases. In the global economic context, one might wish to establish firewalls that could rise up to block the effects of crisis in one market (say, Iceland or Greece or, more dangerous, China or the United States) from spreading to another (any of the other two hundred or so countries in the world). Sometimes the places where interoperability doesn't exist are as important as the places where it does.

The debate about how to get to optimal levels of interoperability too often operates at the extremes. At one extreme, people argue that the state should have no involvement whatsoever in achieving interop. This theory is part and parcel of the dominant strain of cyberlibertarianism, which views any state involvement as anathema to innovation and to the positive development of information and technology systems. At the other end of the spectrum lies the notion that the state ought to drive the development of important complex systems by mandating certain approaches to interop. In this interventionist vision, the leadership of the state is necessary to accomplish high levels of interop; without it, the results will be too uneven and inconsistent to serve the public well.

We, as societies, should not favor one approach or another to interop in the abstract. The type of intervention we choose and who we think should lead it will vary based on a wide range of factors. For instance, when it comes to setting rules for emergency communications, the state ought to be responsible for leading the approach to interop. The state is best positioned to look out for the public's overriding interest in safety and security; also, the type of standard that needs to be set is straightforward. When it comes to determining the best way for e-mail systems to talk to one another, though, there is not much argument against the private sector's lead in terms of setting and managing the standards. In such highly technical cases, the state is ill equipped to make judgments as to standards; that expertise resides primarily in the private sector. Most cases call for a mix of approaches: interop problems tend to be more complex than either of these two simple examples.

To understand the possible options, throughout this book we map a range of approaches that fall along two broad spectrums: private-sector-led approaches ("non-regulatory approaches") versus government-driven measures ("regulatory approaches") on the one hand, and unilateral versus collaborative approaches on the other.

The chart lists the most important interop tools that we have identified in the course of our research into a broad range of examples in the information and communications industry. One way to accomplish interoper-

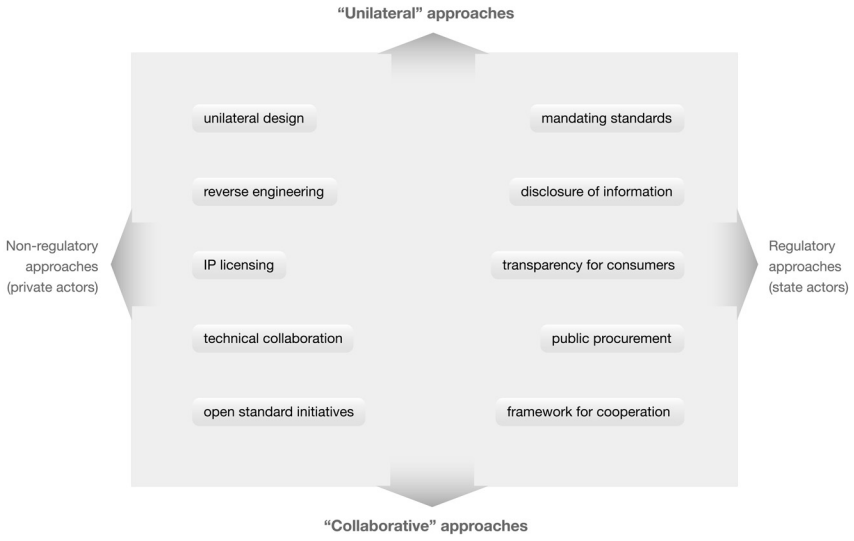

ability is to work within a single firm to interconnect the products that are offered to customers. For instance, Microsoft works hard to ensure that its Word and Excel programs integrate nicely with the Outlook e-mail program and PowerPoint. More often than not, interoperability is accomplished through collaboration between or among two or more firms. Microsoft, for example, has invested heavily in work with Novell to make the two firms' corporate technologies work better together than they used to. In the information business in particular, interoperability is often produced through standards processes, in which interested parties agree to definitions or requirements. They include a wide variety of approaches, ranging from "open" standards processes—that is, open, formal processes administered by standards organizations—to ad hoc cooperation. Office documents are rendered compatible over multiple types of systems, such as Word and OpenOffice, because most firms adhere to open standards for document formats.

In several cases that we have studied in depth—for instance, the business of health care records and the smart grid—market forces alone have not (yet) led to the level of interoperability that is desirable from a public policy perspective. In such instances, governments play a key role in fostering

interoperability. States have a broad range of tools available that help establish or maintain interoperability. Some of these instruments are more invasive than others. Government-imposed standards are a radical form of state intervention. In safety-related areas—such as national security, emergency communications, or navigation—society has an interest in establishing and ensuring interoperability instantly and across the board. For instance, the government mandates that those who sail or drive boats must use certain modes and language to communicate with one another. In other sectors, such as health care or energy provision, government can deploy other tools, incentives, and "softer" approaches. For instance, the state might use its procurement or convening power to induce market actors, such as those who sell health information systems, to aim for higher levels of interoperability.

The state will always be involved to one degree or another simply by virtue of its role in shaping a business environment, legal framework, and regulatory system that can facilitate (or thwart) interoperability efforts across sectors. Nonetheless, there are varying degrees of government involvement, and in turn of private-sector leadership, in the promotion of interop that will make sense. One or more of these approaches used in combination might work to achieve the most advantageous interoperability within a complex system. In most of the cases we have studied, blended approaches—involving diverse actors and one or more approaches concurrently—were applied to increase interoperability. Standards-setting initiatives among private actors that have been facilitated by government agencies are one such example.

It is not enough to achieve interoperability for existing systems. From the design process through implementation, the goal must be sustainable interoperability, to guarantee that the systems will continue to work together. At the same time, it is important to ensure that interoperability over time does not lead to *lock-in*, a situation in which existing forms of interoperability become so standardized that they hamper innovation.

The most informal approaches, such as ad hoc collaboration among firms, are usually the quickest route to interoperability. This can be seen

in the context of the social web: your Facebook account can easily connect to your Twitter account, or you can move a document from Google docs into another part of the Internet cloud with ease, even though these systems were all created by different companies.

It is quite another matter to ensure that you will be able to do so five years from now. Think of all those pictures you've uploaded to Facebook, Shutterfly, Photobucket, Flickr, Picasa, or Kodak Gallery (once upon a time called Ofoto). How do you know you'll be able to get them out, a generation from now, to show your grandkids? How do you know you'll be able to download them? Can you be sure that the data formats will be the same so that you can still view all those photos? How do you know those businesses will even exist? This problem, as we will see, has huge consequences for libraries and for our system of preserving knowledge and information in general. Interoperability can be the solution to these problems over time, too, but only if it is done right.

Interop can serve both to promote innovation and to thwart it. The vexing problem of lock-in hovers at the core of most interop debates. If the system remains flexible in the right ways and at the right levels of the interop stack, then higher levels of interop tend to lead to continued innovation over time. But too much interop, or the wrong kinds of interop, can have the opposite effect, causing a highly interconnected system, such as the global system of air traffic control, to become locked in to the technology of a particular era. We will return many times to the vexing problem of lock-in throughout this book. The lock-in problem helps clarify interop theory as a whole: interop is certainly desirable, but not all the time and not to the highest possible degree in every case.

Our theory of interop establishes a framework but not a single prescription, leaving most of the specifics of how to bring interop about to be determined on a case-by-case basis. That can feel unsatisfying. But it is an essential truth: the most interesting interop problems relate to society's most complex and most fundamental systems. Their answers are never simple to come by, nor are they easy to implement. This characteristic of

interop theory is a feature, not a bug. It is the power of interop both as a lens and as a design principle that is relevant to so many big, intractable, interesting problems. The price to be paid for striving for a universal principle at the level of theory is that such a theory is full of nuances when it comes to application and practice.

We, as societies, must take interop seriously as we hurtle into a future full of increasingly complex and interconnected systems. Interop does not simply help us understand and navigate an increasingly interconnected world; it is also the invisible force that has enabled many great innovations, ranging from transportation systems to the Internet, and it will enable many more. The role of interoperability will become even more important in the future. The responses to the biggest challenges we face as societies, whether climate change or the health care crisis, require the smart use of technologies that connect unimaginably broad sources of information and knowledge, people, organizations, and governments. A sound theory of interop—the art and science of working together—will help break down the barriers that separate us, without creating new problems as we develop the complex systems of the future.