Introduction

IMAGINE FOR A MOMENT that you are lost in New York City without a cell phone or any other way to contact a friend whom you were planning to meet that very same day. Expecting to coordinate at the last minute, you failed to specify a meeting place in advance. You might think it absurd to suppose that the two of you—lacking any way to communicate and lost in the middle of several million people—will ever find a way to meet up. But if you had to pick a time and location in the hope that your friend might be waiting for you at that same place and hour, where would you go, and when?

A common practice would be to wait beneath the clock tower at the information booth in the center of the Main Concourse in Grand Central Station. Is there something particularly suitable about that clock—or even that station—that makes the choice obvious? Certainly, Grand Central Station is well known and, at least since its restoration, very beautiful. It is centrally located, and tourists and commuters routinely pass through it. Perhaps other reasons, too, could be adduced for its attractiveness.

Yet while all of these reasons may matter, none of them matters decisively. None of them makes that particular location a uniquely compelling spot for an unplanned rendezvous, particularly in a city filled with possible places to meet. What does is the established expectation that the clock in the middle of the Main Concourse simply *is* the default place to

meet a friend. And at what time of day would you head to Grand Central? If you are like most people, you would wait under that clock tower at 12 noon—and you would likely find your friend waiting there for you at the same time.

This story about an unplanned meeting in Grand Central Station is adapted from an earlier account that economist Thomas Schelling used to illustrate his idea of "focal points," which are points of reference that coordinate expectations in the absence of prior agreement. Schelling asked an "unscientific sample of respondents" to name a time and location at which they would attempt to meet a friend in New York City without having specified the details in advance and without any way to communicate. An absolute majority of Schelling's interviewees responded that they would go to the clock tower at the information booth in Grand Central Station, and nearly all of them said they would do so at 12 noon. Schelling used this (and similar examples) to illustrate the "tacit coordination" through which, in the absence of express agreement, we nevertheless find ourselves able to coordinate our activities.[1]

Throughout the world, billions of people are similarly looking for places to "meet"—either literally or figuratively—often without having specified the details in advance. Across the globe we ask ourselves a question whose answer in New York City is that clock in the center of Grand Central Station: how should I best position myself in order to "meet up" with other people without a prior agreement? If that omnibus term "globalization" captures anything—and I argue in this book that it does capture something important about our contemporary circumstances—then what it highlights are the diverse but increasingly shared answers to that question.

Globalization involves a game of social coordination similar to that of meeting a lost friend in New York City, except that we are not usually deciding on a location but rather on the languages, laws, technologies, and frames of reference—or, as I refer to them in this book, the *standards*—by which we can best facilitate our newly global activities. We are not so much asking *where* to meet, as *how*. Of course, the "meeting" that these standards facilitate is more complex than locating a friend, but the logic of tacit social coordination is common to both.

For the most part, we are not yet at a stage where the global standards we will use have become clearly known. But in every area of global

activity, as some standards gain prominence, alternative ones become less attractive choices for social coordination. This process can prove self-reinforcing, with the result that a single standard can become the established choice, the convention on which we settle to coordinate global access. Globalization is, among other things, the uneven process by which such conventions are determined, the way in which we construct (or, in many cases, simply receive) the settled terms of access to each other that make international cooperation possible.

The word "globalization" has become impossible to escape and yet remains difficult to define. Indeed, the term now functions in a great deal of scholarship and commentary as a residual category: since almost any contemporary phenomenon of importance crosses some kind of border, the word has become a catchall. Even many serious studies of globalization rapidly degenerate into simple analyses of immediately identifiable global institutions and actors, with little inquiry into their deeper interrelationships or the logic that underlies them. Thus, perhaps *the* defining characteristic of our era receives only piecemeal theorization across various academic disciplines, and little beyond platitudes from public commentators.

GLOBALIZATION AND NETWORK POWER

In this book, I present an argument about how we should understand globalization, claiming that many contemporary phenomena now loosely grouped under this rubric can helpfully be viewed through a single lens. Prominent elements of globalization can be understood as the rise to dominance of shared forms of social coordination, and these global conventions can prove difficult to alter once in place. In areas as diverse as trade, media, legal procedures, industrial control, and perhaps even forms of thought, we are witnessing the emergence of international standards that enable us to coordinate our actions on a global scale. What we are experiencing now, in "globalization," is the creation of an international in-group that welcomes the entire globe on settled terms: a new world order in which we clamor for connection to one another using standards that are offered up for universal use. Yet, while we may all come to share these new global standards—to the extent, at least, that we desire access to the activities that they mediate—we may not all have much influence over their establishment in the first place.

The standards that enable such global coordination display what I call *network power*.[2] The notion of network power consists in the joining of two ideas: first, that coordinating standards are more valuable when greater numbers of people use them, and second, that this dynamic—which I describe as a form of power—can lead to the progressive elimination of the alternatives over which otherwise free choice can effectively be exercised. It is support for, and criticism of, both of these elements, in various guises, combinations, and degrees of self-consciousness, that fuels contemporary debates over globalization.

Network power emerges with the possibility of social coordination via new global standards, made possible by the compression of space and imagination that technological advances have brought. At both the global and the local level, coordination is based largely on expectations. Consider again the problem of meeting your lost friend in New York City. The decision to head to Grand Central Station (and at 12 noon) is determined by a series of reciprocal expectations: where you think your friend will go, which depends upon where he thinks you will go, which depends upon where you think that he thinks you think he will go, and so on.[3] Globally, we face a similar problem: how we are all to "meet" in the global landscape which is now opening up before us requires making best guesses about reciprocal expectations. As any one of the possible solutions to a coordination game becomes a point of reference reflected in these expectations, it generates a form of power, with the capacity to pull in people who might otherwise rely on other conventions.

Although I introduce the idea of network power in the context of a discussion of contemporary globalization, I do not mean to suggest that it is a new phenomenon, but simply that it is one that has become more visible in the contemporary world. Networks, even global ones, are not new—and neither is the power present in the social interactions that generate them. Human history plentifully records intercultural trade, communication, and migration, spanning continents and millennia. But what *is* new about our age is the accelerated emergence of, and linkages among, these global networks. From trade to communication to domestic regulations, what was once mainly, even exclusively, "local" is becoming increasingly global. More precisely, certain versions of local practices, routines, and symbols are being catapulted onto a global stage and offered as a means by which we can gain access to one another. They have become the standards by

which we render each other's actions comprehensible and comparable, and through which we are enabled to engage in forms of beneficial cooperation, such as the exchange of goods or ideas.

Emerging global standards solve our problems—or at least compete to do so—and those standards that gain the most prominence become *focal points* in which we find the solution to the problem of global coordination. This "solution" is, of course, double-edged: it offers coordination among diverse participants but it does so by elevating one solution above others and threatening the elimination of alternative solutions to the same problem. Inherent in the use of any standard is a tension between the cooperation that it allows users to enjoy and the check on innovation that it also imposes, since innovation would constitute a break in an ongoing cooperative regime.

This double-edged aspect allows us to make sense of a seeming paradox at the heart of debates over globalization. On the one hand, globalization is often celebrated as an advance of human freedom in which individuals become ever more able to lead lives of their own choosing. Transnational flows of money, goods, and ideas accompany an increasingly liberal international order in which (it is claimed) individuals are ever freer to participate in a global economy and culture. At the same time, the complaint that globalization is based on power has become widespread, and is developed especially pointedly in recent accusations of "empire." At the center of these seemingly contradictory claims lies a difficulty in untangling voluntary choice-making from coercion, with allegations of both frequently being attached to the same action. For example, the choices of people to learn English or of nations to join the World Trade Organization (WTO) may seem, on the face of it, well reasoned, freely made choices. Yet it is also argued that these choices stem from a kind of domination that they may in turn reinforce, reflecting the systemic power of already privileged actors or institutions. The idea of network power allows us to maintain our common-sense view of people as reasonable, choosing agents while simultaneously allowing that those doing the choosing may be subject to a form of external compulsion. It does so by treating the context of a choice as part of any adequate description of it. On this account, convergence on a set of common global standards is driven by an accretion of individual choices that can be considered both free and unfree. Of course, the claim that our choices can manifest freedom,

in that they express formal consent, and yet still reflect oppression in their recapitulation of systemic unfreedom, is relatively commonplace in modern social theory. It is, however, a claim largely missing from today's discussion of globalization.

Contemporary commentators too often treat the emergence of globalization as either a rational act of global self-construction or an agentless process in which standards spread like free-floating viruses across the planet. In the former case, human agents are conceived as global institution-builders; in the latter, as the hapless victims of a brave new world beyond their control. The idea of network power rejects both of these descriptions as inadequate, but incorporates certain aspects of each, allowing us to see that convergence on a set of global standards can occur purely as the result of the accretion of decentralized choices—choices that can feel both free and forced at the same time.

GLOBALIZATION AS EMPIRE?

Many people greet the arrival of global standards with ambivalence, not only because local practices often get edged out, but because the standards gaining global prominence are not the products of common deliberation and collective effort, but seem to emanate from and privilege certain countries, particularly the United States. One of the complaints about globalization that Americans understand least is that it represents a new empire, or works to the hegemonic advantage of the United States. Some of this incomprehension comes from parochialism and inattention. But it may also seem puzzling that, by all appearances, many of the fiercest critics of globalization either participate in, or aspire to, a global and even "Americanized" way of life. How, one might ask, can such people complain that their free choices are also a form of oppression? These malcontents, it is often said, act in confusion or bad faith, perhaps both.

Network power starts from a different premise, taking seriously the charge that billions live as subjects to a power they feel but whose nature they may not be able to articulate clearly. On this account, the globalization to which we are now witness cannot be described as a straightforwardly voluntary process. Neither can resentment against it be said to represent merely envious discontent. It is rather a challenging response to the global expansion of a form of power that may not have been characterized adequately but which is increasingly dominant in our social relations.

Indeed, the assertion that globalization is imperial has lately become the subject of mainstream and conservative discussion in the United States and elsewhere; it is no longer a charge made by anti-globalization activists alone.[4] The characterization rests on two claims: first, the fact that many important choices (even those taken at the level of whole nations or peoples) can seem to have been already decided on account of the effects of globalization; and second, that many already privileged countries benefit disproportionately from processes of globalization. Both parts of this allegation of imperialism become easier to understand if we consider them in the light of the network power of dominant global standards. Coordination is both liberating and entrapping: liberating because it offers greater access to others, and entrapping because it does so—often necessarily—in a way that privileges one mode of access rather than another.

Of course, the characterization of a form of entrapping coordination as "imperial" may seem exaggerated, given that the term is usually invoked to describe a situation in which one society openly dominates another. The most obvious examples of this control are the political conquest and occupation of foreign societies—as in many cases of the "new imperialism" of the nineteenth century—but the control of one society by another does not require such direct means.[5] The very word "empire" comes from the Latin word *imperium,* which was used to describe the mixture of territorial conquest, informal commercial domination, and cultural hegemony characterizing Rome's rule of the Mediterranean in the early phase of its expansion.

The distinction between formal and informal empire is now a familiar one, serving to contrast a situation in which direct political control is necessary to secure benefits from a subordinated society with one in which it is not.[6] In the latter case, the subordinate society cooperates willingly in serving the interests of the controlling society, whether through relations of economic dependency, military alliance, or some other form of "indirect" control. In any case, the control of one society by another is the same, but the way in which such control is maintained varies. Certainly, as many studies of imperial history conclude, formal and informal imperial strategies should not be thought of as opposed, but may be pursued simultaneously in different geographic contexts or at different moments in time.

The claim that contemporary globalization represents a new kind

of empire rests on the concept of informal empire (or hegemony), since direct imperial control is absent in most of the world. But this idea of informal empire—while intuitively appropriate and often persuasive—leaves unarticulated the precise mechanism of control. Part of the problem is conceptual. To each idea of empire is necessarily tied a model of the power that underlies the control of the subordinate society. Formal domination suggests a model of power operating as the command of a political superior, and which is backed up by outright force. In certain regions of the world, this kind of analysis will seem more plausible than in others. As an account of globalization overall, however, a focus on the political control of territory will fail to offer insight into the economic, cultural and institutional aspects of globalization that are perhaps most interesting —albeit difficult—to capture.

The insight that the relations that seem to liberate us may also bind us emerges clearly in the work of certain theorists of modernity such as Max Weber. But unlike the forms of unfreedom associated with an internally generated modernity, global standards often come (or appear to come) from the outside. They also impose their costs unevenly, and frequently privilege the already powerful. Therefore globalization may appear to many who feel its effects most acutely not as the iron cage of modernity manifest on a newly global scale, but as foreign imposition in the familiar mold of empire. Given these dynamics, we should not expect the accusation of empire to disappear any time soon. But neither should we pretend that it consists of nothing more than confusion and bad faith in a moment of unambiguous global advance. Globalization is Janus-faced, generating new forms of freedom and new problems of entanglement.

TWO FORMS OF POWER

Any account of globalization as empire—or indeed any that sees it as constituted, at least in part, by relations of domination—must confront the problems inherent in theorizing any power relations that do not resemble the command of a political superior over a subordinate. Perhaps the most important of these problems is one that recurred repeatedly in debates in social thought during the nineteenth and twentieth centuries: the difficulty of describing how power can operate through social structures rather than as the express will of a well-defined agent. The concept of network power may have something to contribute to this debate. It can

help us to see how individual actions can create structures that, in turn, limit individual agency in a way that resembles the more familiar exercise of power by one person over another.

Here we can usefully draw a distinction between two dynamics of power that correspond to two different ways in which social activity occurs. One way in which we organize our social life is through the construction of a political will, which allows us to make decisions collectively. Final decisions might still be taken by only one person, by several, or by the entire collectivity, but the process itself can still be described as collective decision-making since it follows a procedure which the entire group accepts as producing decisions that are valid for everyone. The dynamic of power that operates when we take decisions collectively in this way—that is, through a political procedure—can be described as reflecting relations of *sovereignty*. A sovereign decision can reshape social outcomes directly through the exercise of a form of power that resembles the command of a political superior. In the case of democratic government, for example, this occurs when a majority determines an outcome for everyone.

The second dynamic of power operates through what we might call relations of *sociability*. In this case, aggregate outcomes emerge not from an act of collective decision-making, but through the accumulation of decentralized, individual decisions that, taken together, nonetheless conduce to a circumstance that affects the entire group. Market activity and linguistic evolution are paradigmatic instances of the construction of a collective outcome via relations of sociability. This cumulative social construction may—because it rests on voluntary individual cooperation—appear to represent a form of uncoerced participatory activity. While this may sometimes be the case, I argue that such relations of sociability do not generally represent a form of social activity free from the exercise of power. However, the power at work in sociability is not well understood. The idea of network power offers a way to conceptualize the power at work in these relations of sociability, which has today become most visible in the emergence of new global networks. The distinction between sociability and sovereignty has been developed in a variety of important social and political theories, but the distinction between the ways that power operates in each remains underdeveloped. Without a clearer grasp of both of these dynamics, and particularly the differences between them, we cannot properly understand the processes driving contemporary globalization.

OUTLINE OF THE BOOK

In this book, I develop the idea of network power by first offering a general formulation of the concept in broad terms and then analyzing examples drawn from contemporary globalization. However, the book does not proceed by first developing a theory and then offering "applications." Rather, a single argument advances, initially through the abstract presentation of network dynamics, and later in the examination of concrete instances of global networks. It should also be noted that I do not attempt to offer a survey of globalization (or of the ever expanding literature on the subject) but merely select those problems and themes that seem most relevant to my purposes.

In the first chapter, I introduce the concept of network power, the power that a successful standard possesses when it enables cooperation among members of a network. In the way I use the terms, a *network* is united by a *standard,* which is a shared norm or practice that facilitates cooperation among members of a network. Since the reason we use standards is to gain access to others for the sake of cooperation, the more people who adopt a given standard, the more valuable it will be for others to adopt the same one. For example, a language is more valuable for us to learn if many other people already speak it. In general, therefore, the larger the network, the more powerful the standard underlying it will be—and the more pressure non-users will feel to adopt that standard. In extreme cases, this pressure can represent a kind of compulsion to join a dominant network. It is this form of power that generates much of the resentment against globalization, since it may entail a loss of local autonomy and the suppression of alternative modes of coordination.

A standard that possesses network power can develop into a universal standard. A universal standard is a convention on which all members of a network have settled. It reflects a particular kind of cooperative behavior emerging from an interdependence of choice, in which people seek a way to coordinate their actions with one another. The idea of network power attempts to capture the dynamic behind this universalization—the way in which interdependent choices can tend to favor the emergence of a single standard.

As a standard develops toward conventionality, it follows a particular trajectory. This trajectory involves three factors: *reason, force,* and *chance*. The path of any real-world network is, of course, highly contingent upon

historical trends and may well involve all three factors in combination. Studying the trajectory along which standards move reveals that the power of any particular standard may be unrelated to the inherent benefits it has to offer. For example, the English language is today poised to become the global language in many domains not because of anything intrinsically special about it, but because of Britain's imperial history and the success of American mass media. After an initial push that may come from a variety of causes, the network power of a standard alone may be sufficient to drive it toward conventionality.

In Chapter 2, I locate the idea of network power in relation to the distinction between *sociability* and *sovereignty,* and then set that distinction in the context of the history of social thought more generally. The idea of network power may be understood as an attempt to analyze the dynamics of the power operative in sociability: that is, as a way of understanding how power is structured through social relations outside the formal politics of sovereignty. An analysis of this sort is particularly important in order to characterize globalization adequately. Indeed, we can understand the process of globalization as one in which the relations of sociability tend to outstrip the relations of sovereignty, the latter being contained within the boundaries of nation-states in a way that the former are not. Such an understanding also allows us to think clearly about individual agency and social structure together in the context of contemporary globalization, neither exalting individual agency while neglecting the possibility of entrapment in social structures, nor elucidating social structures without considering how individual agents can transform them.

In Chapter 3, I make these arguments more concrete by examining the dominance of the English language today and the dominance of the gold standard in the nineteenth and early twentieth centuries. Both of these cases provide interesting examples of social coordination via shared standards. Language allows communication, and money facilitates economic exchange. The symbolic tokens that enable linguistic and monetary coordination possess network power, and English and gold are particularly interesting examples of standards that have attained (or did once attain) global dominance. Studying the contemporary rise of English shows us a standard as it progresses toward universality. Examining the rise and fall of the gold standard reveals the tensions generated when a universal standard is unable to accommodate the need for local autonomy and innovation.

As a given standard becomes dominant and moves to universality, it eclipses rival standards that formerly facilitated the same activity. It may also prevent the emergence of alternative forms of coordination in the future. The choice to adopt the dominant standard then becomes an increasingly coerced one, for the only options are to join it or face social isolation. It is the prevalence of this kind of choice—a circumstance in which I argue a kind of power is at work—that makes the emergence of global networks something other than a large-scale act of international voluntarism, or the free enactment of a global social contract. Accordingly, a standard may be selected not for its intrinsic properties but for the access it provides to the network it coordinates—or even, in extreme cases, on account of the concomitant risk of losing access to any social network whatsoever. A dominant standard may even prove intrinsically disadvantageous in terms of the activity which it coordinates, and yet possess enough network power to compel non-users to adopt it.

In Chapter 4, I explore the idea of network power *as power* more systematically, examining different conceptions of power and situating the idea of network power within these ongoing debates. An understanding of the distinction between sociability and sovereignty is present (if only implicitly) in many of the most important social theories of the nineteenth and twentieth centuries, and especially in attempts to develop heterodox theories of power. For instance, the idea of *hegemony* presents a model of power in which consent and domination are not opposed, but in which people who are subjected to domination consent in some way to the power exercised over them. Yet ideas of power based on hegemony or related concepts have had a mixed reception, in part because they can be evocative but vague, depriving us of a clear way of thinking about human agency and social structure together. I suggest that the idea of network power can provide a theory of power in which consent and domination are also joined, but in which the relationship between human agency and social structure is better elucidated.

Of course, arguing that something constitutes "power" says little about whether that power is normatively justified. We cannot be *against* power independent of an evaluation of the uses to which it is put, or the legitimacy of the circumstances that give rise to it. In Chapter 5, I outline the elements of a normative assessment of network power, focusing on two general classes of concerns: first, a concern about the extent to which

different people are able to realize their interests, and second, a concern about the maintenance of cultural identity. With regard to the former, I discuss theories of distributive justice, and to the latter, theories of recognition, in order to consider how an inequality of network power might be evaluated in these frameworks.

Where we judge an instance of network power abusive, how best are we to counter it? In Chapter 6, I turn to possible strategies for defusing network power where we judge it harmful. It is difficult to determine which strategies to deploy against network power because, unlike in cases of straightforward coercion, this kind of power is driven by consent. Against many forms of power we rely on the elaboration of rights—both *negative rights* that guarantee a zone of individual autonomy and *positive rights* that guarantee a minimum level of resources. I argue that the provision of negative rights will fail to remedy abuses of network power, precisely because network power is consent-driven. The strategy of elaborating positive rights also leaves much to be desired since it is difficult to specify how the state (or some other party) might intervene in support of a threatened standard. Instead, we need to reconceive both kinds of rights relationally, asking who we think ought to be free, from what constraints or because of what enabling conditions, to undertake which actions.

Adopting a relational approach along these lines, I propose that the best way to counter network power is through institutional changes to the configuration of networks. Since we are clamoring for access to one another, the only way to manage network power is to provide alternative and multiple channels for such access, refusing to privilege any single one. The question of whether this is possible takes us far from the idea of rights, whether positive or negative, and suggests instead a more contextual, institutional analysis that understands the extent of our effective liberty to depend, not on any abstract conception of right, but on the details of the institutions in which we make our lives. Accordingly, defusing network power in any concrete instance requires that we first understand how it inheres in different network configurations. I examine three "network properties"—availability, compatibility, and malleability—and suggest how different combinations of these properties affect the power of a given network. Understanding at a deeper level of specificity what kinds of standards gain network power may enable us to change the institutional context in which that power arises.

In Chapter 7, I examine the spread of global technologies, including technical standards and practices. The most striking technological development of the last few decades has been the emergence of telecommunications technology and personal computing on a global scale. These technologies are dependent on underlying standards for technical coordination and therefore present striking examples of the dynamics of network power in action. I discuss the anti-trust case *United States v. Microsoft,* and its implications for the regulation of network power in high technology. I then turn to the "open-source" movement prominent in new forms of collaborative digital production. I argue that these egalitarian relations of production can be preserved only by embracing, rather than rejecting, politics. I therefore criticize the anarchistic tendencies of many adherents of the open-source movement, arguing that only if the democratic politics of sovereignty are mobilized on its behalf will a free digital world be able to survive.

In Chapter 8, I analyze the network power of the World Trade Organization. To some, the WTO is a major achievement of the international community: a victory of progressive forces over parochialism and unenlightened self-interest. To others, it is a telling example of a world driven to undemocratic centralization by economic forces beyond the control of ordinary people. Understanding the WTO as a standard that exerts network power will make these divergent views intelligible and allow us deeper insight into the organization and its possible future. The debate over the WTO is obviously related to the debate over trade liberalization more generally, so I also take this opportunity to look more deeply into arguments in favor of such liberalization and the ends it serves, in order to explicate further the network power dynamics at work in this domain.

The WTO is the most visible institutional embodiment of neoliberalism or, as it is sometimes referred to, the "Washington Consensus." Neoliberalism serves as a portmanteau term for a set of economic policy proposals that privilege markets and criticize collective interventions in the economy. In Chapter 9, I discuss these policies and argue that neoliberalism is not one standard but a host of related standards, the network properties of which are not generally configured to provide them with network power. Neoliberal economic policies sometimes govern access to others—to the benefits of foreign aid, foreign markets, and foreign recognition—and when they do they may coordinate access via a stan-

dard that possesses network power. At other times, however, countries adopt neoliberal policies for reasons unrelated to network power—for example, because a particular policy seems persuasive to a country's decision-makers, or because of direct pressure from a foreign country or international agency. Because neoliberal policies possess varying degrees of network power, the Washington Consensus is not a "take it or leave it" standard. Rather, it can be broken down, adopted piecemeal, and revised, as many successful economies have managed to do. The way in which neoliberal policies can become coordinating standards that do possess network power is through a process that I term *juridification*—the creation of formal legal principles governing a mode of access, as in the example of the WTO. I explore the process of juridification in the negotiations over a global agreement governing international investment, which produced the dramatic and controversial failure of the Multilateral Agreement on Investment (MAI).

One of the widespread complaints against globalization is that it produces, or hastens, cultural loss and homogenization, moving us toward a dystopic "McWorld." Accordingly, in Chapter 10 I turn to the subject of cultural globalization. Unlike specific economic policies or technological platforms in which the governing standard and the power it exerts may be relatively clear, this form of globalization is governed by complex standards in which the dynamics of network power are much less obvious. Of course, the constituent elements of "culture" are very diverse; I adopt the framework introduced by the sociologist Peter Berger in order to examine the "four faces of global culture" in the light of network power. I then explore what are sometimes called "epistemic communities," networks of people with shared expertise, such as the international community of scientists. If we can understand the shared conceptual frames linking these communities as standards, then network power may generate cultural convergence, not just at the level of consumer preferences but even in forms of thought.

The rise of universal standards such as the ones I discuss in this book link us together in global networks that are often in tension with many of our important "local" commitments. This is most significantly and, in my view, worryingly the case in the mismatch between the global scope of these relations of sociability and the domestic locus of democratic sovereignty. It is clear that the globalization of social relations cannot

be straightforwardly balanced by a concomitant globalization of the structures of political authority. But whether a globalized sociability will mollify or exacerbate antagonisms among different sovereign states is a long-standing question to which history offers unsettling answers.

For it is not the first time that the world has been in this situation. John Maynard Keynes's famous description of the world before the "Great War" may seem eerily familiar to contemporary eyes:

> The inhabitant of London could order by telephone, sipping his morning tea in bed, the various products of the whole earth, in such quantity as he might see fit, and reasonably expect their early delivery upon his doorstep; he could at the same moment and by the same means adventure his wealth in the natural resources and new enterprises of any quarter of the world, and share, without exertion or even trouble, in their prospective fruits and advantages . . . [he] could then proceed abroad to foreign quarters without knowledge of their religion, language or customs, bearing coined wealth upon his person, and would consider himself greatly aggrieved and much surprised at the least interference. But, most important of all, he regarded this state of affairs as normal, certain and permanent, except in the direction of further improvement, and any deviation from it as aberrant, scandalous and avoidable.[7]

If we are to avoid the bloodshed that ended this earlier episode of globalization, we must press forward to a global modernity that we can live with. This is a task for politics above all, but our political strategies may be limited by our conceptual analysis. Our first step, then, is coming to terms with globalization as it is now unfolding.

CHAPTER ONE

Defining Network Power

WITH THE END OF THE COLD WAR, a world divided into hostile halves suddenly became "one world" in the middle of a historic transformation: the integration and consolidation of activities not just at the national or even continental level, but on the global scale too. Commerce, technology, media, and cultural imagery spilled across national borders in what appeared to some commentators to be a new worldwide free-for-all. Thus was the revolution of "globalization" suddenly upon us—as it has been, in fact, for the past few centuries.

In most contemporary discussions, globalization is presented as constituting a break from the past, a contemporary circumstance without precedent. But as the historian Emma Rothschild reminds us, not only does the idea of globalization have a history—even the idea of globalization as a phenomenon *without a history* has a history.[1] The failure to grasp the long history of globalization epitomizes an ongoing failure to understand its real significance today. Part of the problem is that we lack a plausible framework in which to understand the diverse but arguably interrelated phenomena that constitute "globalization," past or present. In this chapter, I suggest a framework for making sense of these diverse phenomena, as they have unfolded historically and as they appear in the world today.

THEORIES OF GLOBALIZATION

The social processes that we group together under the rubric "globalization" stretch back to the early modern period, and often before. Writers in Western Europe—which has sometimes been a fairly isolated corner of the world—have been commenting since at least the sixteenth century on the increasing internationalization of their commercial, intellectual, and cultural affairs, and on the changes in the subjective identities of people and peoples that such internationalization brought about.[2] These early commentators on what we now call globalization were aware that technological and social changes had led to a compression of space (as it might be put today) and that this compression had myriad consequences for their commerce, their political and social life, and their self-understanding. The French writer and diplomat François-René de Chateaubriand made an argument that sounds strikingly contemporary when he wrote, in 1841, that technological advances could be expected to bring about an international society: "When steam power will be perfected, when, together with telegraphy and railways, it will have made distances disappear, it will not only be commodities which travel, but also ideas which will have wings. When fiscal and commercial barriers will have been abolished between different states, as they have already been between the provinces of the same state; when different countries, in daily relations, tend toward the unity of peoples, how will you be able to revive the old mode of separation?"[3]

It is a rather parochial conceit of contemporary commentators, therefore, that globalization is something unique to our time. But while the processes of globalization are not new, its theorization under this name is of more recent provenance—and this fact helps to account for the claim that globalization is itself novel. The word "globalization" in its current usage goes back only as far as the 1960s, while even the related word "international" only originated (with Jeremy Bentham) in the late eighteenth century.[4] Sustained analyses of globalization are of even more recent vintage, most fewer than two or three decades old.

Any idea of globalization—whether in Chateaubriand's time or in the present—must begin with the compression of space, a change in geographic distance as it is lived and conceived. At a subjective level, at least, this change is undeniable. The widespread feeling that "the world is getting smaller" attests to the fact that technological and social advances have made the distance between points—and people—on the globe feel

far less significant than they did only a short time ago. Objectively, the distances *are* shorter, at least if we think of the time required to traverse them. With modern transport, no two cities in the world are, any longer, any more than about a day's travel apart. The world no longer seems to hold points unfathomably far away, geographically or culturally.

The sociologist Anthony Giddens, an able contemporary commentator on this feature of globalization, offers a much-cited definition in these terms: "Globalisation can . . . be defined as the intensification of worldwide social relations which link distant localities in such a way that local happenings are shaped by events occurring many miles away and vice versa. This is a dialectical process. . . . Local transformation is as much a part of globalisation as the lateral extension of social connections across time and space."[5]

Giddens understands globalization as a process through which distant localities become linked in a way that constitutes a "lateral extension of social connections across time and space." This process also transforms the nature of these localities in relation to the global. In a similar way, one of the first scholars to discuss "globalization" under that name, Roland Robertson, emphasizes this compression of distance: "Globalization as a concept refers both to the compression of the world and the intensification of consciousness of the world as a whole . . . both concrete global interdependence and consciousness of the global whole."[6] Drawing on both these descriptions, the social theorist Malcolm Waters has defined globalization as a "social process in which the constraints of geography on social and cultural arrangements recede and in which people become increasingly aware that they are receding."[7] These scholars recognize that our experience of global geography has changed, that we are collectively aware of this change, and that it alters many "local" aspects of our lives that have now come to possess a new "global" dimension.

The idea of network power originates with this observation of geographic compression, but treats it only as a prelude. In its most essential aspects, globalization occurs *after* this change in the experience of global geography. The release from geographic constraint does not in itself bring about many transformations in the way that people live and relate to each other without a second set of social changes that enable them to cooperate in international and transnational activities. Chateaubriand writes that when "distances disappear," then "ideas have wings." But, of course, ideas

do not really have wings: it is people who promote particular ideas, and who come into contact with other people who may at some point come to share them. It is not technological advances themselves that "tend toward the unity of peoples," but the fact that social relations can now be conducted on a global level, among greater masses of people, thanks in part to those advances.

What the compression of geography enables but does not provide is social coordination. That coordination, I argue, is achieved in the adoption of shared *standards* that allow global social networks to emerge following the technological changes that bring people into contact with one another, but do not (in themselves) generate everything else that people require to become mutually intelligible and beneficial partners in cooperation. This social coordination is achieved via standards that provide the frameworks for global cooperation—the languages, points of reference, customs, rules, laws and regulations—that must follow the compression of space for the creation of a global society.[8]

On this account, globalization is the disruptive and uneven process by which we come to share common standards after the eclipse of distance. Because there are usually multiple possible standards that might enable global cooperation in any specific domain, the choice of any one of them is socially and politically consequential. Unsurprisingly, then, the struggle over these standards—over the forms that our emerging globalization could take—is one of the defining features of our age.

STANDARDS AND NETWORKS

In this book, I advance an argument about how we might understand the social coordination that standards provide, focusing on the power that standards have in bringing into being new global networks. The idea of a "network" has come into greater usage in recent discussions of globalization. However, since the concept is used differently in different scholarly disciplines, let me be explicit about how I will use the term. A *network* is an interconnected group of people linked to one another in a way that makes them capable of beneficial cooperation, which can take various forms, including the exchange of goods and ideas. While every network is ultimately composed of people, the way in which these people gain access to one another may include mediation, for example through computers or corporations. Networks can also exist in many different

forms with important implications for network dynamics, a point to which I will return in Chapter 6.[9]

A *standard* defines the particular way in which a group of people is interconnected in a network. It is the shared norm or practice that enables network members to gain access to one another, facilitating their cooperation.[10] A standard must be shared among members of the network to a sufficient degree that they can achieve forms of reciprocity, exchange, or collective effort. Consider, for example, networks of English speakers, Internet chat room participants, consumers in the Euro zone, or people who use the metric system. In every case, a standard is central to the existence of the network, serving as a convention common to all its members. Without it, we would see a collection of isolated individuals rather than a connected group capable of achieving something together.

A network can be, but is not necessarily the same thing as, a "community," at least if we mean by that term a geographically proximate and socially integrated group. Global networks may well not function or feel like communities to those participating in them. When we imagine a community, we tend to envision a close-knit group of people who share a great deal—values, language, and, usually, some specific geographic location. By contrast, a network may be rather "thin," united in one particular fashion and perhaps in that fashion only, linking participants who may be scattered physically across the globe. At its most diffuse, networks of computer users around the world "meet" over a network of wires and satellite signals constituting a virtual network of people who have never actually met one another personally—nor are ever likely to do so. But global networks need not be technologically based; they may result from the common use of a particular currency, language, regulatory regime, or other convention.[11]

It is important to differentiate between two kinds of standards: what I call "mediating" standards and "membership" standards. These two types of standards are not mutually exclusive, and distinguishing between them can be considerably complex in any given instance.[12] Both of these kinds of standards can exert network power: mediating standards do so inherently while membership standards need to be configured in a particular manner in order to do so, as I will explain further in Chapter 6.

What I mean by a *mediating standard* is a standard that governs access to others by its very nature; some particular social activity is inherently

regulated by it. Language is perhaps the most obvious example of a mediating standard: to join the network of English speakers, you must learn English. Mediating standards are ones that cannot be avoided if users wish to engage in certain activities: they form a part of that very activity itself. Thus they serve as solutions to the problem of social coordination, thereby enabling beneficial cooperation in a network.

We also use the term "standard" in a second sense, to refer to a specific ideal, exemplar, or required level of attainment, as in discussions of a "standard" in fashion, hygiene, education, environmental protection, or industrial manufacture. This second kind of standard is not necessarily basic or inherent to a given activity (as the English language is to communication with people who only speak English) but rather establishes an ideal or target. Thus, these standards do not generally govern the access that we can have to others, except where membership in a network is predicated upon the common acceptance of such a standard as a precondition for access. When it is, we can identify it as a membership *standard*. Membership standards do not enable beneficial cooperation by serving as coordinating standards, but they do govern access to an in-group by specifying criteria for admission to a network.

Mediating standards, because they inhere in a given social activity, are self-enforcing standards. No one needs to insist on the use of English in London or U.S. dollars in Florida; the logic of decentralized social coordination is inherently regulative in these cases. By contrast, membership standards usually require enforcement by some actor or set of actors to exclude all but those who adopt particular norms. For example, membership in the European Union (EU) is predicated on the acceptance of the *acquis communautaire,* which is the minimum set of regulations required of all EU states. (The *acquis* consists of the entire body of laws of the EU, including all treaties, regulations, directives passed by the EU institutions, and the judgments of the European Court of Justice.) The *acquis* serves as the membership standard regulating the entry of new countries to the EU, each of which is required to "adopt, implement, and enforce" it upon its accession.[13]

STANDARDS AND ECONOMIES OF SCALE

This description of standards may give the impression that they are relatively unremarkable, little more than social facts the general existence of which

we all take for granted. Standards exist everywhere, and we have a sense of how they work: by making things comparable, commensurable, comprehensible, and thus accessible and available. But the most interesting feature of standards—and the one that gives rise to what I call "network power"—is that they can "spread," propelled by people's desire for access to members of a network. A commonly used standard uniting a desirable network will be attractive to any outsiders wanting to gain access to that network and such outsiders may well decide to adopt it for that purpose alone.

To illustrate this point, suppose for a moment that you are the head of a manufacturing company, *Bolt Inc.*, which produces small mechanical parts, such as bolts. In the United States, bolts are identified by two measurements: their diameter in fractions of an inch, and the number of threads per inch. Now imagine that a large conglomeration of the major producers and corporate consumers of mechanical parts calling itself the "Alliance for Rational Standards" commits itself to switching over from the imperial standard of measurement to the metric standard. Metric bolts are described by their diameter and "pitch"—the distance between two threads—given in millimeters. Members of the Alliance claim that the adoption of the metric standard will bring them into conformity with mainstream global practice, and bridge the gap between the production of precision scientific instruments, which are always designed in metric, and that of retail products sold to Americans, which are designed according to the imperial standard. A credible announcement of this kind will force other producers and consumers of small mechanical parts to make a decision about the standards they will use. For Bolt Inc., the question now is whether or not to follow the Alliance and switch from the imperial to the metric standard, manufacturing bolts in thicknesses measured not in fractions of an inch, but in millimeters. Making this switch may require reconfiguring or replacing existing bolt manufacturing machines, which we will assume can only produce bolts according to one specification and so would not be able to produce in both metric and imperial measures. But your overriding concern will not, in all likelihood, be with the potential mechanical difficulties of switching over. Your principal consideration will be your expectations about Bolt Inc.'s customers, who own machines designed to take a particular kind of bolt, built to imperial specifications. If these customers go metric and Bolt Inc. does not, Bolt Inc. will lose their business. Will Bolt Inc. go metric?

Assuming that the market share that the Alliance represents is significant—and further, that you expect that other non-Alliance companies in the market will also want to switch to metric—you will almost surely adopt the metric standard too. Otherwise, you risk losing a large part of your business: that is, all the clients for whom your bolts will no longer prove compatible once they have switched, as well as any future clients who switch over at a later point. Of course, this might mean losing some clients who prefer the imperial measurements that you currently use, but your considered choice might necessarily involve such a loss. It might also mean accepting the cost of changing some equipment. But if the Alliance's total market share passed a critical threshold, you would probably be forced to make the switch to metric anyway, in order to continue to compete. Fortunately, your clients would presumably do so as well: in fact, they may decide that they are at risk of losing Bolt Inc. as their supplier of bolts if they do not join the Alliance, even in advance of any decision you make.

The power of standards generates these reciprocally determining decisions in the market for bolts. Borrowing a concept from economics, we might say that there are *economies of scale* in the adoption of metric, as proposed by the Alliance. (The notion of economies of scale is formally distinct from, but intuitively similar to, the more familiar notion of "increasing returns to scale."[14]) A production process exhibits such economies of scale if there is a decreasing cost for each additional unit of output over some range of production. In this case, it is more efficient to produce on a larger scale since each additional unit of production costs less than the previous unit did.

The value of a telephone offers an example of economies of scale. How much is a telephone worth? One answer might be that a phone is worth what it costs to produce, the sum total of its parts plus the labor needed to assemble them. However, we do not buy the telephone for the machine itself, but for the *connection* that it offers to others who also have telephones. So how much it is worth depends on how much we want it, which depends in turn on the number of people who already have a telephone (or whom we expect to get one). If the world contained only one single telephone, it would be worth nothing to us—even though it costs something to produce—for there would be nobody in the world we could call. Given the existence of other telephone users in the world, however,

our phone begins to look more valuable. And when billions of other people own telephones, that same telephone, far from being worthless, becomes extremely valuable—so much so that many people might even be prone to claim that they "could not live" without one.

The economies of scale that we see in the telephone network are also exhibited in the example of Bolt Inc. The greater the number of people who use a particular standard, the more valuable it becomes for others to adopt the same one. Expectations play a critical role in this process. Even if the Alliance does not currently control a large market share, its announcement may well make large numbers of people believe that the metric standard will become widely used—leading many of them to adopt it preemptively. Expecting that its customers (and its competitors) will go metric, Bolt Inc. will probably decide do so. In turn, this will confirm the expectations of other market participants, who, anticipating Bolt Inc.'s switchover, will likely have done so too. Generalizing this dynamic, we can easily imagine a scenario in which, before long, there may be no companies left which are still using the imperial standard: most likely every company will have either made the switch or gone out of business. In this example, we see how decentralized, individual decisions to participate in the adoption of a particular standard can conduce to a circumstance in which the standard becomes effectively compulsory for everyone in a given market—and all without any collective decision to produce such an outcome ever having been taken.

THE NETWORK POWER OF STANDARDS

The economies of scale driving the adoption of a standard result from a "positive feedback" dynamic in which each new user increases the desirability of that standard in the eyes of other potential users. A system is said to exhibit positive feedback when a change in one variable leads to a further change in that same variable, and in the same direction. For example, an increase in the number of people who own a telephone increases the value of the telephone network as a whole, which then attracts new telephone users. Or, the greater number of people who speak a given language, say, the more attractive that language is likely to appear to those who wish to learn a foreign tongue. It is the positive feedback generated by the adoption of a standard that constitutes its power, the "pull" that a standard has because it underlies a network of users to whom others

are likely to want to gain access. The concept of *network power* joins two ideas: first, that standards are more valuable when greater numbers of people use them because they offer a form of coordination that exhibits economies of scale; and second, that one effect of this coordination is, over time, to eliminate alternative standards that might have been freely chosen. When these ideas are considered together, the central premise of network power is that the benefits that come from using one standard rather than another increase with the number of users, such that dominant standards can edge out rival ones. This process exhibits a positive feedback dynamic and can prove self-reinforcing: the value of a standard increases with the addition of each new user to its network, which means that it has the power to draw in additional users from other networks, each of which further increases its network power. Importantly, when a user switches from one network to another, she will not only increase the value of the network she joins, but will also decrease the value of her previous network by leaving it with one fewer member (in turn, making it less attractive for the remaining members of that standard to continue using it, and more attractive for them to copy her defection).

Although there is a sort of social "momentum" behind network power, the growth of a network is driven by the active *choices* of individuals, rather than by their passive acceptance of something external to them. An analogy to physical momentum (describing network power as a "snowball effect") can tend to obscure the agency of the people involved, and direct attention away from the decisions of network members and toward properties considered at the aggregate level. But the concept of network power should not tempt us to locate active agency in aggregate social phenomena such as "the network" rather than in the people who compose it. Networks do not grow of their own volition, or "spread" in an impersonal way, even though it is easy to drift into employing that kind of language when we want to describe how network power leads to growth in network size, and how that, in turn, increases network power. It is *our* choices that lie behind network power, and nothing beyond or outside them—a point I will return to in the discussion of modern social thought in the next chapter.[15]

It is easy to see network power at work in the technological examples I introduced above. But in many forms of social exchange, too, a relatively specific standard underlies a network and regulates access to the benefits of cooperation with members of that network.[16] The network power of

these standards is a measure of the attractiveness of that network to outsiders. Since it is via such standards that people gain access to one another in networks, whenever it is desirable to participate in a larger network rather than a smaller one, *standards have a power that grows in proportion to the size of the network they unite*.

In computing, the idea that the value of a network is related to its size has been given technical elaboration in "Metcalfe's Law," which asserts that the value of a computer network increases exponentially with the number of its users. However, the broader point can be made more generally: standards attract new users in proportion to the size of the network they underlie, and sometimes in greater proportion.[17] Of course, the fact that a standard has network power does not mean that non-users will necessarily adopt it. A standard is attractive as a gatekeeper to a network of users, even if that network is not large or otherwise significant enough to induce non-users to switch onto that standard. Importantly, network power is always a *comparative* notion, based on the different sizes of rival networks—that is, networks based on different standards, each of which facilitates the same activity. Even the standard underlying a small network generates network power, for any standard that mediates access to others is valuable. Whether it is valuable enough to adopt is another question.[18]

The size of the network that a standard unites is not the only factor that can lead one standard to dominate over others. A network may be attractive to outsiders not only because of its size but because of the desirability of the particular standard that unites it—a point to which I will return below. Also, the idea of "size" requires further elaboration, for it is always a comparative measure. It refers to the proportion of people engaged in a particular activity who use one standard rather than another, not the absolute number of people using a standard. For example, if there were only ten people in the world participating in a certain activity, seven by one standard and three by another, the seven-person network would be considered a "large" one. Networks may also vary in significance not on account of their relative size, but on account of the significance of their members in the activity that the standard mediates: a small handful of elite financiers may have much more effect on the use of a particular currency than numerically larger networks of currency users in poor countries who command fewer resources. The network power of standards that mediate

especially important or desirable activities will also be greater than that of those that mediate less central or important activities. The benefits of being part of the club, so to speak, are greater in some fields of endeavor than others.

SWITCHING NETWORKS

Network power can induce people to "switch" networks—to do things like learn foreign languages, use different currencies, or join organizations that require adherence to new rules of conduct, all of which occur frequently in contemporary globalization. To illustrate more precisely how network power can lead to such changes, consider in the abstract two standards, Standard A and Standard B, which unite their respective networks, Network A and Network B. These standards enable the same activity, and function in roughly equivalent ways: that is, the difficulty of the activity does not change depending on which standard is used. People in Network A would gain access to everyone in Network B if they used Standard B. Likewise, members of Network B would need to use Standard A to deal with people in Network A. Individual members of either network will, whenever it is desirable to gain access to a greater rather than a smaller number of people, encounter strong incentives to use both standards or, where they must use only one or the other, to use the one that provides access to the larger network.

For example, suppose that Network A has many more members than Network B. Any member of either network would want to be able to use both Standards A and B if they are compatible, thereby gaining access to all the members of both networks. But if these standards are incompatible, such that they govern the same mode of social interaction without allowing complementary or parallel structures—an assumption I will re-examine in Chapter 6—then network power will push members of Network B to adopt Standard A and join Network A. The incentive to do so will come from the possibility of gaining access to the larger number of members of Network A, even at the cost of losing access to the remaining members of Network B.

Examining the incentives faced by an individual member of Network B makes the point clearer. The member of Network B will need to consider the benefits of cooperating with the larger body of members in Network A that she would gain if she adopted Standard A. These benefits would be

offset by whatever costs she would incur in the process of switching over from Standard B to Standard A, which might include losing access to the remaining members of Network B and the difficulty of adjusting to the new standard. (We might call the first set of costs "opportunity costs," and the second "switching costs," but remember that staying put in Network B also generates its own set of "opportunity costs"—the loss of potential cooperation with members of Network A.) So any member of Network B contemplating the move into Network A will have to compare the value of accessing Network A against the costs of adopting Standard A and losing the value of current membership in Network B. Since the value of a network is related to its size, we can conclude that when Network A is larger than Network B, any member of B will face pressure to join A.

A greater disparity in the size of the two networks will translate into a greater discrepancy in network power. This does not mean that Standard B has *no* network power. Because cooperation with members of Network B is desirable, Standard B will exert some pull on members of Network A. But this pull must be assessed comparatively. All else being equal, users of Standard A will remain in the dominant Network A, sacrificing potential access to members of Network B in order to maintain a connection to the larger network.

So far, we have been assuming that Standards A and B are equivalent in their functioning, so that the choice between Networks A and B is based on their value as a function of their respective sizes only. In abstract models of networks (or in some actual computer networks) this assumption may hold, since the connections of which the network is made up may be considered uniform. But in the case of standards that enable social interaction, we cannot assume that standards operate similarly in different networks. Rather, the value of membership in such networks must be considered a function of both the size of that network and the quality of its standard.

We may think of this difference as one of *intrinsic* and *extrinsic* reason. We have *intrinsic* reasons to adopt a new standard if it simply functions better for our purposes, while we have *extrinsic* reasons if that standard governs access to a network that we find desirable for some other reason, most likely because of its size. Thus the decision to switch from Network B to Network A, where Standards A and B perform differently and are of differing quality, will depend on the size of the two networks *and* the

different levels of performance of the standards in comparison, an analysis that includes both intrinsic and extrinsic reasons. I will discuss extrinsic and intrinsic reasons at greater length below.

So far, assuming a picture of the two networks at one moment in time—a snapshot that shows Network A to be larger than Network B—the costs and benefits facing a member of Network B considering the switch to Standard A will look as I have described. But taking a dynamic approach that examines how networks grow and decline over time—and incorporates the expectations of such growth or decline—reveals the positive feedback effect which generates network power. Given the economies of scale in the adoption of a standard, members of Network B who do not switch to Network A will face *increasing,* rather than constant, costs for staying put.

Suppose that members of Network B, the smaller network, face costs of different magnitudes in switching over to Standard A. Very plausibly, the adoption of a new standard will prove difficult in varying degrees for different people, depending on a whole host of factors, including their capacity for learning, their attachment to and position in the original network, and so on. We might imagine the members of Network B distributed along a line illustrating the costs to each member of switching over: from near zero, where we put those to whom the cost of switching to Standard A will be least, to some large number at the other end, where we put those to whom the cost of switching over will be greatest. Those with the least cost of switching will presumably do so first, since they face only the opportunity cost of losing access to Network B, and very few switching costs of any magnitude. But this first wave of departures will make the value of Network A increase, precisely because of the increase in its size. Thus the opportunity cost of remaining in Network B will not remain constant. It will increase every time a member of Network B defects to Network A, since with each lost member, Network B will become smaller in comparison to Network A, and hence less attractive to potential members, while Network A will grow in attractiveness. Each new defection from Network B to Network A will likely trigger another wave of defections, drawn from those members of Network B for whom the cost of switching over is now least. As this process continues, the costs incurred by those members of Network B for whom switching over was always more difficult will themselves increase—perhaps even at an in-

creasing rate—as the opportunity cost of staying in the declining network continue to mount up. This means not only that Network A will progressively increase in size (and hence attractiveness) but that Network B will become progressively less viable as a real alternative to A.

Of course, members of a small network will not always wait until mounting costs make switching to the dominant network necessary. Frequently, they will anticipate this dynamic and act in advance of it, expecting waves of defections to occur, and not wanting to be the last to switch over. Members of Network B will rationally expect that, given some minimum necessary level of imbalance between Network A and Network B, Network B will progressively lose members to Network A, and that holdouts in Network B will face steadily mounting costs so long as they do not adopt Standard A. These expectations then become self-confirming and reinforce the power of a dominant standard; especially if the switch comes to be viewed as inevitable, members of Network B can be expected to start scrambling to join Network A in the hope of avoiding being the last to switch over, since any delay in switching will be ever more costly in terms of lost access to the growing Network A. The expectations of members of both networks will determine the expected consequences for any individual member's decision, which will be reciprocally factored into the expectations of all other members in turn.

Understanding the role of expectations in network power helps to explain how the promulgation of standards can be a form of strategic interaction. For example, in the case of the "Alliance for Rational Standards" discussed above, even the expectation that the Alliance's standards will become dominant may push them to becoming so in actual fact. With a credible commitment to these new standards by key companies, the network power of the Alliance's standards will induce non-users to adopt them, even if the dominance of these standards is not already accomplished but merely expected.

INTRINSIC AND EXTRINSIC REASONS

So far, my presentation of the concept of network power has focused on cases of competition between two existing networks, each of which possesses some network power. But in order to clarify the process through which particular standards come to be adopted, we need to take a step back. Whenever a new standard emerges, the reasons for which it is used

will have little to do with network power, since a standard will develop that power only once the network it unites is itself substantial enough to seem attractive to potential users. Leaving network power aside for the moment, then, we can say that there are three main causes for which a particular standard might be adopted: *reason, force,* and *chance*. The origination and development of any actual network will be highly contingent, and may well involve all three causes in different combinations at different times throughout its history. Put differently, we can say that whenever we enter into cooperative arrangements, we will either be pursuing our interests or values (using reason), acting under duress (being subject to force), or doing so by accident (being subject to chance). Reason, force, and chance will continue to act as contributing factors in a network's development, and one way of understanding network power is to describe it as a combination of one kind of reason—extrinsic reason—and one kind of force—indirect force.[19]

By *reason,* I mean simply that one standard is preferred to the alternatives, because it *better meets the goals that an agent has when entering into cooperative relations*. The reason for the preference of one standard over another may be *intrinsic* or *extrinsic*. By intrinsic reasons, I mean that a given standard is attractive for its inherent properties and not because it happens to unite an already large network. For example, certain standards may simply function better, facilitating easier and richer forms of social interaction, or allowing their users to gain access to each other with less difficulty. The claim that one standard is, strictly speaking, superior to another will often be very contentious, requiring a careful delineation of the evaluative framework in question as well as the empirical content of the standards being compared. But it will usually be the case that a standard that enables the desired coordination easily will be intrinsically preferable to more difficult alternatives. For example, in eighteenth-century Sweden, the monetization of copper meant that people had to cart around wheelbarrow loads of coinage (rather than simply carry a small purse of gold) in order to conduct large transactions. For such transactions, a gold standard rather than a copper one would better facilitate economic exchange.

A standard is adopted for *extrinsic reasons* when it is selected not for its internal characteristics, but because of the size of the network it unites. If a greater number of potential partners for cooperation use one standard rather than another, that standard will be more attractive extrinsically. It

is not unreasonable to adopt a standard for the value of the network it unites, but this consideration is rather different from choosing based on its intrinsic properties. Indeed, in many cases, the only reason to adopt one standard rather than another may be extrinsic, particularly if there are no important intrinsic differences between the standards. In the example above, the decision to switch from Network B to the larger Network A was owing to extrinsic reason, determined by the size of the network rather than anything about the underlying standard, since I specified at the outset that the two standards were of roughly equivalent functioning—that is, of roughly equal intrinsic value.

This distinction between intrinsic and extrinsic reasons generalizes a set of insights from linguistics about the nature of words as arbitrary conventions. The Swiss linguist Ferdinand de Saussure based his analysis on the arbitrary nature of the sign. He insisted that the choice of particular sounds or letters to serve as a linguistic sign was wholly arbitrary in that there is no necessary or intrinsic relation between a word and that to which it refers. By arbitrary, Saussure did not mean that signification was open to individual revision, but that there was no intrinsic relationship between what he called a signifier and what it signified; he rejected the possibility "that a signal depends on the free choice of the speaker."[20]

The use of words is both conventional and arbitrary: there is no reason that we should indicate the animal *canis lupus familiaris* by the word "dog" in English and "*chien*" in French—or *canis lupus familiaris,* in biological nomenclature, for that matter—other than that we do so conventionally in order to be understood by other English or French speakers. A given word "has no natural connexion in reality" to that to which it refers; nothing in the sound pattern (or textual elaboration) of a word ties it intrinsically to what it is used to indicate. Thus, in the terms I have introduced here—which follow Saussure's own reference to the possible "intrinsic" qualities of a sign[21]—words have only extrinsic and not intrinsic reasons for their adoption, leaving aside the odd case of onomatopoeia like *bang* or *quack*. Saussure recognized that some signs do bear a connection to their referent, and he called these signs *symbols*. A symbol always retains what Saussure described as "a vestige of natural connexion between the signal and its signification," possessing what he elsewhere identified as a "rational connexion with what it symbolizes."[22] Although he did not discuss symbols at length in his *Course in General Linguistics,* it is possible

to make (as later work in semiology has done) the argument that symbols are conventional, but not entirely arbitrary. Symbols may be taken up for a combination of extrinsic and intrinsic reasons: extrinsic because they are conventional and thus take their ultimate value from common use, and intrinsic because they can nevertheless do a better or worse job representing something.

Saussure clearly recognized the conventional and non-arbitrary value of symbols: "any means of expression accepted in a society rests in principle upon a collective habit or a convention, which comes to the same thing. Signs of politeness, for instance, although often endowed with a certain natural expressiveness . . . are none the less fixed by rule. It is this rule which renders them obligatory, not their intrinsic value."[23] For the sake of expositional clarity, however, Saussure focused his attention on linguistic signs because "signs which are entirely arbitrary convey better than others the ideal semiological process," even while acknowledging that other conventions may be distinguished according to their intrinsic qualities.[24]

Later work in semiotics has distinguished the purely arbitrary conventions of signification (such as words in natural languages) from the iconic mode (in which the sign is supposed to resemble that to which it refers, as in onomatopoeia or in models) and the indexical mode (in which the sign is directly connected to its referent, as in a photograph or a recording). Consider the difference between gold coins and paper money: both are used to represent value, but the former retains a direct connection to something intrinsically valued (bullion) while the latter is a purely arbitrary convention. The insight here is that not everything that is conventional—that is, governed by social conventions—is for that reason arbitrary, lacking any intrinsic connection between the signifier and signified.

My distinction between extrinsic and intrinsic reasons is meant to capture these basic semiotic insights and generalize them beyond language alone. Other than in systems of purely arbitrary signification, such as language as described by Saussure or paper currencies, standards will be selected for a combination of intrinsic and extrinsic reasons. But as the network power of a standard grows, the intrinsic reasons why it should be adopted become less important relative to the extrinsic benefits of coordination that the standard can provide: the conventional value of a standard will come to outweigh any intrinsic merits or demerits.

The British adoption of the metric system offers a good example of the role of extrinsic reason in the choice among (non-linguistic) conventions, since the switch from imperial to metric measurements was made in order to allow smoother exchange with Britain's European neighbors. Certainly, we can imagine intrinsic reasons for the adoption of metric, such as the ease of calculating with a decimal system. But the extrinsic coordination—the fact that the metric standard was used in the rest of Europe—proved decisive. Indeed, even if the metric system had provided an intrinsically *less* reasonable standard, Britain would still probably have adopted it for the sake of easier coordination with its trading partners. By the same token, had the opposite had been the case and if the metric system had been used only in Britain, with imperial measurements used throughout the rest of Europe, Britain would most likely have abandoned the metric system, regardless of its intrinsic advantages.

Over the course of the rise of a standard, the importance of extrinsic reasons in deciding whether or not to adopt that standard increases relative to the importance of the relevant intrinsic reasons. The initial growth of a network may be based on intrinsic reasons, but as that network grows over time, the attraction of its size will provide a sufficient extrinsic reason for its adoption. This does not mean that an extrinsic reason displaces any intrinsic reason for choosing a standard, but simply that the matter becomes overdetermined. The intrinsic merits of a standard will usually remain relevant even with the increased effects of network power, but they will be supplemented—and perhaps, eventually, wholly outweighed—by extrinsic reasons. For example, it may have been *intrinsically* reasonable for Britain to adopt the metric system at any point in the history of the decimal-base measurement standard, whether in the year 1800, following the French Revolution that introduced it, in the year 1900, at the height of the British Empire, or in the year 2000, with European economic integration the order of the day. The intrinsic merits of metric have not changed over the last two hundred years, but the weight of extrinsic reason bearing down on the British did—and that is what proved decisive.

FORCE AND CHANCE

Of course, the early growth of a network does not usually depend on the intrinsic reasons for adopting one standard rather than another; history seldom presents such a straightforward choice. Another major factor in

the early adoption of one standard rather than another is *force*. By force, I mean the *use of coercion* to compel the adoption of a standard. Force may be either direct or indirect: direct force should be understood as the imposition of costs unrelated to network membership that compel a switch, while indirect force is the imposition of (opportunity) costs resulting from membership in one network rather than another.

Direct force involves imposing "costs" such as violence or punishment for failing to adopt a given standard, or denying benefits unrelated to the immediate standard-governed activity. For example, in the category of direct force, we can include not only threats or acts of physical violence but also coercive socialization, such as education in which pupils are not allowed to speak to one another in their native language, or are punished for retaining traditional dress or customs. If direct force is the motivation for switching networks, then we can say that a person adopting a new standard has been compelled to do so outright and not because of anything about the configuration of networks. In such cases, the person making the switch is forced to do so by the threat of punishment or deprivation unrelated to her membership in the network, although this factor will, of course, have an impact on any decisions she has to make about her network membership.

Indirect force, by contrast, is the pressure to adopt a standard that comes from the threat of losing access to others, the social isolation from people who use a different standard. Indirect force comes not from the imposition of costs unrelated to network membership, but from the imposition of costs driven by the structure of network membership itself. The opportunity costs that I described above as mounting on those who hold out against a dominant standard should be understood as a form of indirect force. A member of one network faces overwhelming indirect force when these opportunity costs mount to such an extent that she is deprived of real opportunities for cooperation in any network except the dominant one. But a lesser degree of indirect force may also be felt in any case of the rise of one standard over another. Importantly, the penalties extant in indirect force consist not just of lost access to users of a dominant standard by users of a smaller one but, in the case of great inequality or disproportion in network power, lost access to former users of one's own standard who have defected to the larger network as well.

Early in the rise of a network, only direct rather than indirect force

may be used to increase network size; if force is to be used, new adherents must be added by direct coercion, since indirect force emerges only once a network of sufficient size exists that non-users feel compelled to join it or face social isolation. The crux of the formulation of network power *as power*—a theme I discuss in Chapter 4 at length—is that beyond a certain point, members of a small network are "forced" to adopt the standard of a dominant network or else face isolation. For example, an individual's decision to adopt a majority language (when it is not driven by direct compulsion) comes from the desire to communicate with other speakers of that language, but also because a minority language becomes progressively less viable as the population of its speakers dwindles. To be a part of a viable network *at all* may require defecting to a larger one.

On this account, indirect force can be seen as equivalent to extrinsic reason, for the same benefits of coordination that draw in new users based on extrinsic considerations also generate the structure of opportunity costs that I describe here as indirect force. The attraction of a standard underlying a dominant network constitutes both a form of indirect force and an extrinsic reason, either or both of which can be used to describe choices made because of network power. Thus, early on in the rise of a network, we may see either intrinsic reasons or direct force augmenting its membership, but after the network passes a certain size, both extrinsic reasons and indirect force will come into play, since the standard that unites it will have become valuable for the social coordination it provides.

Indeed, indirect force and extrinsic reason may be said to merge at the highest levels of network power, for the demand of reason is to adopt a dominant standard rather than lose access to others, and yet this reason dictates a selection that, while rational, is chosen under the compulsion of having no viable alternatives. When reason and force merge under conditions of great disparities in network power, it becomes a matter of semantics whether we should best describe the situation as determined by "reason" or determined by "force." Recognizing the equivalence of indirect force and extrinsic reason enables us to see both at work in the same process: intrinsic reason becomes irrelevant and direct force unnecessary once a standard has surpassed a critical threshold and its continued rise is driven by network power.

Finally, a standard may attract early adherents merely by *chance,* by which I mean the *accidental convergence on a routine*. We may accidentally

converge upon a standard that provides the solution to a problem of coordination, whose resolution in any one of many possible ways is desirable. And once we have stumbled upon such a solution, even if we suppose it is not the best solution possible, it may well continue to have some staying power. Consider in this regard the arbitrary nature of spoken words: if there is no intrinsic connection between the phonetic qualities of a word and its referent, then we may suppose its emergence to have been an accident. (Therefore, Saussure proposed studying signs synchronically rather than diachronically, rejecting the historical approach of nineteenth-century comparative grammar and arguing instead that words have meaning as they are used in relation to a system of signification at a given moment in time.[25]) Other standards, too, may emerge as the result of happenstance or fortunate convergence, as in the case of "pure conventions," which will be discussed in the next chapter. The literature on industrial organization is filled with stories of products that gained market dominance because of seemingly insignificant differences in initial market share that later became magnified because of economies of scale. In such circumstances, an early lead gained by accident can generate unexpectedly large differences in final outcomes. However, although we may fall into the use of one standard rather than another by chance, we cannot always adopt a new standard just as easily. While chance may be responsible for the initial emergence of a standard, it will be transformed into network power—into indirect force and extrinsic reason—as the standard gains new users. For, once on the trajectory of network power, it is not accident that will increase the size of a network, but its own amassed power. Put differently, even arbitrary signs possess network power, precisely because of their conventionality.

THRESHOLDS IN NETWORK POWER

After a standard gains its initial push from reason, force, or chance, it exerts network power and moves along to conventionality, buoyed by an influx of new users. We have examined how this process begins, but two further thresholds on this trajectory remain to be delineated in order to complete the trajectory of a standard. Like "tipping point" phenomena, a standard can surpass critical thresholds in the path of network power, exhibiting new properties after gaining critical mass.[26] Importantly, these thresholds are subjective and not necessarily determinate or objective;

I introduce them in order to capture the social and phenomenological evaluation of successful standards.

The first threshold is what I call the *threshold of visibility,* signifying the point past which a network is sufficiently large to become attractive to non-users. Below this threshold, a standard has no network power because it is effectively invisible to those outside its network. Network power arises beyond this threshold, once the standard has attracted a large enough number of users that its network becomes visible and attractive to others. Reason, force, or chance (or some combination thereof) may push a standard past this threshold. Beyond it, the network begins generating network power, exerting a pull on members of other networks for extrinsic reasons. Below it—though the network exists—it is not large enough even to register with outsiders.

Language offers a convenient way of exemplifying what I mean by a threshold of visibility. Many English speakers would like to learn French or Spanish, two languages with large populations of speakers both in France and Spain and in their former colonies abroad. Both French and Spanish have passed the threshold of visibility and prove attractive to English speakers. (Whether they prove attractive enough for English speakers to make the effort to learn them is another question.) But the language of the Eyak, now known only to the elderly Marie Smith of Cordova, Alaska, remains below the threshold of visibility—so far below that it seems certain that with Marie's death it will be extinguished altogether as a living language.[27]

As this example illustrates, standards do not only rise above the threshold of visibility; they can fall below it, too. Indeed, when smaller networks decline, whether through the death of their members or their members' absorption into larger networks, these smaller networks can fall below the threshold of visibility and become "invisible," powerless in the broader sphere of social relations.[28] These standards may still have users but they have no network power: total effacement of a standard is not necessary for it to lose all effective relevance to non-users. Not only Eyak, but a very large number of dying languages have fallen below this threshold, as I discuss in Chapter 3. A small community may still speak them, but they are effectively invisible to outsiders, even in a very local geographic setting.

Below the threshold of visibility, standards do not exhibit economies

of scale in their adoption; they do not possess network power. Hence, the maintenance of such a standard may suffer many of the problems of non-conventional cooperative behavior: for example, familiar problems of collective action. Individual defection from the standard may be individually preferable to cooperation, even if cooperation is in everyone's collective interest. (However, it should be noted that this suggestion depends, in part, on a view of "free riding" that may be philosophically untenable.[29]) It is perhaps unsurprising, then, that when standards are deliberately created—work that may involve large initial investments—the effort will often have been made by powerful players with strategic ambitions and a long-term planning horizon, such as large companies, government agencies, or industry consortia. The move beyond the threshold of visibility requires such a subsidy. Network power will take over beyond it, but cannot be counted on before the network attains some minimum size.

The second threshold is what I call the *threshold of inevitability,* signifying the point past which a network has become so dominant that we can expect virtually all non-users to adopt its standard. Convergence on the dominant network has become inevitable; this threshold comes before convergence on a universal standard, but after its imminent arrival has come to be expected by everyone. This threshold is a kind of tipping point, as we see in studies of non-linear dynamic systems, beyond which the diversity of possible alternatives is quickly distilled into a single unified network. We may understand this threshold as occurring at the point at which the expected value of joining the dominant network to all remaining non-users (as a function of both its size and the quality of its standard) is equal to that of any lesser network plus switching costs. Remember that by this point, the value of a smaller network to its members may be very small indeed, as its total size will have decreased due to prior defections, and its current value will reflect a general expectation of further attrition.

A critical question to consider is whether one network's movement beyond the threshold of inevitability will tend to push competing networks below the threshold of visibility. As a general matter, the movement of one standard past the threshold of inevitability will tend to displace competing standards to the extent that they prove incompatible with it.[30] In the extreme case, if two standards are *strictly* incompatible—such that a user can use only one at a time—then as the dominant standard moves toward universality, any rival standards will necessarily fall below the threshold of

visibility and may even be pushed to extinction, for a universal standard that is strictly incompatible with its rivals will be (by definition) the only game in town. But the movement of one standard to universal status does not require the elimination of compatible alternatives, only incompatible ones. For example, the rise of English as a global language should prove compatible with the continued existence of other languages, if it is used as a global *second* language rather than the sole global language. However, in other cases—for example, some technical or monetary conventions—the rise of a universal standard may lead to the effective elimination of alternatives that prove incompatible.

Again, these thresholds may be identified through phenomenological or social assessment: they are not necessarily distinct points, and they are certainly not invariant across different standards. Rather, they represent loose concepts whose actual content will vary given the empirical circumstances of any particular network. There is probably no abstract formula by which we can determine the network size at which either tipping point occurs, unlike in some models of non-linear systems. For example, whether the threshold of inevitability is passed after a network gains a precise minimum size or is a more gradual process occurring once a network is recognized as clearly dominant cannot be addressed out of context. Similarly, whether the threshold of visibility is as small as two people using a standard with each other in a purely private communication, or as great as a large minority of a given population is also an empirical, not an abstract matter, and depends on the relevant network configuration in question. But these thresholds capture two defining aspects of network power: first, that a network requires a minimum size to exert power over members of other networks, and second, that beyond a certain size, a network pulls all other networks with it toward convergence on a universal convention.

THE UNIVERSAL STANDARD

It should be obvious from the discussion so far that network power will push different networks toward a position of convergence, in which two or more networks become one. When all members of Network B have switched over to Network A—or at least, so many that Network B no longer possesses any attraction for members of Network A—then we can say that Network B has converged with Network A. If we continue the assumption

of strict incompatibility, then we can say that Network B has fallen below the threshold of visibility and Network A has surpassed the threshold of inevitability and successfully universalized.[31]

It is important to note that convergence on the same standard should not be assumed to be the only outcome possible when two networks come into contact—various other positions of equilibrium are possible, and this particular one depends on several initial assumptions, particularly of the strict incompatibility of the standards in question and their equal intrinsic value. It may not be the result, for example, if the intrinsic value of two standards differs greatly, or is perceived to differ greatly by two relatively distinct groups of people who wish to maintain separate networks. Whether someone will choose to join a large network with an inferior standard depends on a balance of intrinsic and extrinsic reasons that need not tend toward convergence.[32] In the general case, however, network power will tend to drive networks to convergence on the same standard, because standards are inherently universalizing—even though the domain in which their predominance grows may be restricted in any number of ways. The social relations embodied in cooperative arrangements like standards suggest themselves as incipient universals. Every coordinative arrangement represents an invitation to others because it proclaims a manner in which social cooperation can be accomplished.

We can call a standard whose invitation has been taken up by everyone a *universal standard*.[33] The power of a standard thus follows a peculiar trajectory on its way to universality: starting from reason, force, or chance, it grows in relation to the size of its network—even at an increasing rate—up to the point at which it replaces all competing standards. At that point, when a network is as large as it can be without becoming the only viable network, the network power of its standard is at its zenith. Beyond that point, however, we cannot properly speak of the standard retaining any power at all.[34] Rather, its power has become so total that the comparative notion of network power is no longer applicable, for all its rivals have fallen below the threshold of visibility. Such a standard admits no others against which we can compare it. It has become conventional: a commonplace, a universal.

The triumph of a universal standard is revealed in both its generality and its obscurity. It is difficult to identify a universal standard in the absence of contrasting alternatives. Put differently, we might say that

conventions are most conspicuous when partial. A universal standard may often prove inconspicuous, appearing to its users simply as a social fact with its actual function in social coordination obscured. We notice social facts only with great difficulty for they constitute our modes of relation to others and influence our perception and interpretation of these modes themselves.[35]

The great tumult of globalization, past and present, results from the great diversity of standards brought into competition with one another and thus rendered visible in a new way. And our attention is primed to that diversity by the fact that, almost as soon as it is registered, it appears to be at risk of melting away in a series of global convergences. Understanding why this may be occurring requires examining more closely the network power at work in these newly global social relations.