Journal of Theoretical Politics 19(1): 33–54
DOI: 10.1177/0951629807071016
http://jtp.sagepub.com

CONCEPTUALIZING CONTINUITY AND CHANGE

THE COMPOSITE-STANDARD MODEL OF PATH DEPENDENCE

Taylor C. Boas

ABSTRACT

Political scientists studying institutional development face the challenge of accounting for both continuity and change over time. Models of path dependence based on increasing returns, inspired by the example of the QWERTY typewriter keyboard, have played an important role in the analysis of institutional continuity, but they have been criticized for their inability to accommodate change. In this article I present an alternative model of path dependence inspired by the example of the Internet, a technology that has changed fundamentally since its invention. The *composite-standard* model of path dependence illustrates how complex political institutions subject to increasing returns can evolve gradually over time through a changing mix of lower-level component parts. By incorporating mechanisms of institutional change, such as conversion and layering, within an increasing returns-based theoretical framework, the composite-standard model highlights new interconnections among these previously distinct processes and offers new insights into the nature of long-term political change.

KEY WORDS • historical institutionalism • increasing returns • Internet • path dependence • QWERTY

Introduction

Scholars who seek to advance generalizable models of institutional development face the basic challenge of accounting for both continuity and change over time. Indeed, as Thelen (2003) has argued, we are often drawn to the study of institutions precisely because they change dramatically in certain ways while remaining remarkably stable in others. Building upon models of technological development

I am particularly grateful to David Collier for multiple rounds of comments on earlier versions of this article. Others who provided helpful feedback include Ruth Berins Collier, Todd LaPorte, James Mahoney, Helen Oliver, Paul Pierson, Eric Schickler, Steven Weber, John Zysman, two anonymous reviewers, and participants in the Berkeley Latin American Politics Research Seminar, Fall 2003–Spring 2004. The research for this article was supported by a National Science Foundation Graduate Research Fellowship.

first elaborated in economics (David, 1985, 1986; Arthur, 1989), an initial wave of scholarship in political science and sociology applied the concept of path dependence to political institutions, emphasizing lock-in and increasing returns (Pierson, 2000), self-reinforcing sequences (Mahoney, 2000), and the 'mechanisms of reproduction' (Collier and Collier, 1991) of particular historical legacies. These works played the important role of systematizing arguments about historical causation and moving the discussion of institutional stability beyond the arguably unassailable notion that 'history matters'.

More recent research on institutional development has sought to move away from the concept of path dependence, focusing on the ways in which institutions change rather than remain stable over time. Thelen (2003, 2004), Hacker (2004), Crouch and Farrell (2004), Schwartz (2004), and Alexander (2001) have all argued that prevailing models of path dependence overstate the degree of stability in political institutions. In advancing their own explanations for institutional development, the majority of these scholars have distanced themselves from the notions of increasing returns, lock-in, and switching costs that were central to the economics literature on path dependence and also played a key role in Pierson's (2000) and Mahoney's (2000) formulations.[1] Rather, they have focused on new mechanisms of institutional change, including layering and conversion, which they characterize as distinct from the increasing returns to scale that drive traditional models of path dependence (Thelen, 2003, 2004; Hacker, 2004).[2]

In this article I take up the ongoing challenge of elaborating models of institutional development that account for change over time. In contrast to most recent statements on this issue, however, I argue that in order to model institutional change, it is useful to re-examine and build upon the increasing returns-based models of path dependent technological development initially formulated by economic historians. By revising rather than distancing ourselves from path dependence models based on increasing returns, we not only salvage the insights that they have already offered. We also stand to gain new insight into mechanisms of institutional change such as layering and conversion, including the ways in which they are intimately connected with increasing returns, rather than being separate and distinct.

I begin this article with a critical analysis of the technological examples of path dependence in the economics literature, in particular the development of the QWERTY typewriter keyboard. As the empirical example with which the term

1. Crouch and Farrell (2004) represent an exception in this regard, in that their explanations for institutional change build upon Arthur et al.'s (1983) generalized Polya urn scheme, one of the traditional models of path dependence in the economics literature.

2. Hacker (2004) proposes an additional mechanism of gradual change, 'drift,' in which an institution (or policy) remains stable but its overall effect changes because of shifting external circumstances. Because only the *effect* of the institution is transformed, I do not consider drift in this article, which focuses on *institutional* change.

path dependence was first associated, and as a technology with which arguably all scholars are familiar, QWERTY has proven a popular reference not only for economists but also for political scientists seeking to explain patterns of continuity over time. The QWERTY keyboard is resistant to change in part because of the significant barriers to incrementally rearranging the layout of keys. Yet this very feature becomes an important limitation when scholars use models inspired by QWERTY to explain the development of political institutions that may exhibit considerable flexibility, rather than rigidity, over time.

In place of QWERTY, I suggest that better models for many patterns of long-term political change can be derived from technologies such as the Internet, whose standards are *composite* rather than *simple*. Because users of the Internet coordinate around a macro-level standard that is made up of many component parts, the technology itself can evolve through changes in this collection of parts, without reducing the network effects that encourage further adoption. This process of development is still path dependent in that increasing returns discourage abrupt technological shifts at any one moment. Nonetheless, the example of the Internet shows that composite-standard technologies are also capable of dramatic changes in the long run.

The real payoff for political scientists, of course, comes not from explaining technological development *per se*, but rather from using the models that can be derived from these examples to explain patterns of continuity and change in political institutions. When the composite-standard model of path dependence is applied to the study of institutional development, it illustrates how incremental changes in political institutions can cumulate into a fundamental transformation over time, even as increasing returns render an institution resistant to wholesale change at any given moment. In particular, the composite-standard model has the advantage of highlighting the close connections between such mechanisms as increasing returns, conversion, and layering, unifying them under a single theoretical roof. While sacrificing parsimony, composite-standard technologies offer more general and often more accurate models of the complex processes of institutional evolution that historical institutionalists typically examine.

The QWERTY-inspired Model of Path Dependence

While arguments about the historical origins of present-day political institutions have deep roots in the social sciences, from the work of Gerschenkron (1962), Moore (1966), and Lipset and Rokkan (1967) to Stinchcombe's (1968) writings on historical causes, the specific concept of path dependence is a more recent innovation. Political scientists' understanding of path dependence incorporates many insights from these classic studies, but it also builds upon the initial elaboration of the concept in the field of economic history, in particular, Paul David's (1985, 1986) paradigm-defining account of the development of the QWERTY typewriter

keyboard. In this section I examine the history of QWERTY and consider how the features of this empirical example have been incorporated into more general models of path dependence. In particular, I demonstrate that the QWERTY model does not accommodate the possibility of fundamental change over time because the arrangement of letters on the keyboard is resistant to incremental alteration.

The story of the QWERTY typewriter demonstrates how an inflexible standard for keyboard configuration was locked into place through increasing returns to scale.[3] In 1867, Christopher Latham Sholes patented a typewriter design in which the printing action was hidden from view. The major disadvantage of this design was that the user would not necessarily know if a typebar got stuck and would likely continue typing, thus repeatedly hammering the stuck letter into the paper. Sholes sought to perfect his invention by rearranging the layout of the keys; he ultimately settled on the QWERTY configuration, which was less likely to jam because typebars for common letter pairs would strike from different sides of the machine. Other typewriter designs soon became available, including configurations that obviated the problem of typebar jamming, but by then QWERTY had established an early lead in the market.

The QWERTY keyboard's initial popularity was crucial for its ultimate success because typists and typewriter owners each benefited from the market being dominated by a single keyboard configuration. Businesses got greater value if they owned the typewriter on which a majority of typists had been trained, and budding typists wanted to train on the machine that was most common. Once proficient on a QWERTY keyboard, typists were reluctant to relearn the typing process on a new (and possibly more efficient) machine because of the time, effort, and minimal payoff associated with retraining on a typewriter that no one else used. For this reason, argues David, the Dvorak keyboard introduced in the 1930s never stood a chance in the marketplace against the less efficient but universally accepted QWERTY.

David's QWERTY example soon proved to be the foundation for developing more general models of path dependence in economics. Inspired in part by the QWERTY story, both David (1987) and Arthur (1989) developed mathematical models of the adoption of competing technologies subject to increasing returns, and Arthur et al. (1983) also elaborated the model of a generalized Polya urn scheme to capture the dynamics of path dependence in a broad variety of substantive areas.[4] While Liebowitz and Margolis (1990) have since challenged many of the specific historical details of David's account of QWERTY's development, he has continued

3. The following account is drawn largely from David (1985). Liebowitz and Margolis (1990) have contested some of David's historical evidence about QWERTY's inefficiency, but not this basic account of its origins.

4. In a simple Polya urn scheme, a ball is drawn at random from an urn containing an equal proportion of red and white balls. The ball is then returned to the urn along with another ball of the same color, increasing the probability that future draws will be of that color. When repeated, the process eventually settles into a unique equilibrium distribution of red and white balls. This equilibrium distribution depends greatly on the random outcome of the initial draws.

to defend the utility of the typewriter example as 'an instructive and empirically sound heuristic, exhibiting a constellation of generic features to which many episodes in technological and institutional history conform' (David, 1999: 8).

A number of generic features in the history of QWERTY that were incorporated into models of technological development distinguish a path dependent process from the standard neoclassical economic model of markets (Arthur, 1988, 1989). The process is *unpredictable*, characterized by multiple possible equilibria, and it may ultimately achieve an *inefficient* equilibrium due to imperfect information or other sources of market failure. Path dependence is also characterized by *nonergodicity*, meaning that events occurring early in a path are not averaged out and forgotten.[5] Furthermore, path dependent processes result in *lock-in* through increasing returns. The costs of switching to a previously discarded alternative accumulate over time, rendering such wholesale change less and less likely.[6]

The mechanisms sustaining lock-in are important to consider in detail, for they have played a central role in applying the concept of path dependence in political science. Four features of a technology can generate increasing returns to scale (Arthur, 1988, 1989). First, the technology may have a *large ratio of fixed to marginal cost*, so that the production cost per unit declines as production increases. Second, the technology's adoption may be characterized by *learning effects*: the more it is used, the more its efficiency can be improved vis-à-vis other alternatives. Third, path dependent technologies often display *coordination or network effects*, in which the demand for a technology (and its value to each current user) increases with each additional unit sold.[7] Finally, path dependent processes of technological development are characterized by *adaptive expectations*. Users select the technology that they believe will become the most common, and in doing so they increase the probability that others will make the same choice.

The elaboration of models inspired by QWERTY has facilitated the analysis of other examples of path dependence in economics, and it has also encouraged careful scrutiny of the degree to which these models correspond to the real-world

5. In Arthur's (1994: 27) analysis, these first three features depend upon the occurrence of small events that 'lie outside the main description of the dynamic structure' of a path dependent sequence. Such events are not completely random; they are merely exogenous to the causal model that explains path dependence itself. This characterization explains the focus on contingency in recent work on path dependence by political scientists and sociologists. Like Arthur, Mahoney (2000) does not require contingent events to be entirely random, but rather unexplainable by the theory used to account for dynamics within the path dependent sequence.

6. Arthur (1989) has also referred to this final characteristic as inflexibility. I find this term misleading, since flexibility can also be used to refer to the nature of a standard, that is, how much change is possible within a path, rather than how difficult it is to switch completely to another path. In this article, I use flexibility or rigidity to refer to standards, and lock-in to characterize difficulty in switching paths.

7. While often conflated with the effects of high fixed and low marginal costs, network effects are a separate mechanism in that they involve increasing demand for a more widely used technology rather than a lower cost to supply that technology in the marketplace (Lemley and McGowan, 1998).

processes that they seek to describe. On the latter front, Liebowitz and Margolis (1995, 2001) have spearheaded an often polemical but arguably useful debate over whether inferior technologies actually do triumph in the marketplace. In particular, they maintain that the models in which David and Arthur ground their arguments about market failure omit crucial features of the economy such as futures markets and property rights. On this basis, Liebowitz and Margolis (2001) question the applicability of these models in a number of empirical realms, including the competition between VHS and Beta videocassette recorders and between DOS and Macintosh computer operating systems.

While much of the economics debate has focused on this question of efficiency, a less heralded feature of QWERTY-inspired models of path dependence is especially relevant for political science: their emphasis on stability over time.[8] As Leibowitz and Margolis argue (2001), stability is not an inherent characteristic of path dependent processes; on the contrary, it derives from the rigid rather than flexible nature of the technological standard involved. In the empirical example of QWERTY, the stasis-enhancing effect of a rigid standard can be seen quite clearly. Moving even a single key on the QWERTY keyboard would necessitate some degree of retraining and would reduce the value derived from coordination effects. Being locked into QWERTY means being locked into the same arrangement of keys, which has been carried over from typewriters to computers and wireless communication devices. The only realistic potential for change in this case is through an exogenous shock, the wholesale replacement of keyboards themselves by some successor technology such as voice recognition.

The Composite-Standard Model of Path Dependence

In the real-world example of QWERTY, the rigid nature of the technological standard is an empirical fact, but in more general models of path dependence, one need not make the assumption that the standard engendering coordination is incapable of change over time. Indeed, both technologies and institutions often display a capacity for long-term change even as their development follows a trajectory that can meaningfully be described as path dependent. In the search for alternative ways of modeling the development of political institutions, therefore, it is useful to consider the example of technologies whose development displays more flexibility over time than the QWERTY typewriter keyboard.

In this section I develop an alternative model of path dependence – the composite-standard model – which is inspired by the empirical example of the Internet. Like many other communication technologies, the Internet's development displays a

8. The debate over efficiency is less important for political science, since there is no common assumption in the discipline that the natural workings of politics result in efficient outcomes.

distinctly path dependent trajectory: the accumulation of switching costs discourages the wholesale shift to another technology at any given time (David, 2001). Because the Internet's diffusion involves coordination around a composite rather than a simple standard, however, the technology as a whole evinces a degree of flexibility that is not matched by its component parts. Since many of the institutions of interest to historically oriented political scientists are themselves composed of micro-level institutions, the composite-standard model of path dependence provides a basis for reasoning about the juxtaposition of both continuity and change in political institutions. While increasing returns render them resistant to radical change at any one moment, both technological and institutional processes that conform to the composite-standard model of path dependence may involve the accumulation of fundamental changes over time.

From Control Frustrating to Control Facilitating: The Evolution of the Internet

Since its origin in the 1960s, the inherent flexibility of the Internet has allowed for a shift from a fairly anarchic technology to one more amenable to government and corporate control. Initially, the Internet was deliberately designed as a technology that would not lend itself to centralized control. In contrast to the telephone network, in which centralized switching stations route calls from sender to receiver, the 'brains' of the Internet are located in the millions of computers and individual networks attached to it, while the Internet's core protocol, TCP/IP, simply governs the delivery of pre-addressed bits of data to their destinations (Isenberg, 1997).[9] This particular design, known as 'end-to-end' (Saltzer et al., 1984), means that one cannot control the entire network through control of a small number of centralized nodes. Control can be exerted at the ends of the network, but as these ends multiply, controlling the entire network by controlling the ends becomes less and less feasible.

The Internet, however, was not only designed to be resistant to centralized control; it was also designed to be a highly flexible technology, capable of evolving and ultimately thriving in environments that its designers could not possibly have foreseen. At the time of its creation, there was little sense of what services the Internet would need to support in the future, so the core of the network was built as a set of simple, flexible tools. Any service that conforms to TCP/IP can be implemented at the ends of the network without altering the center. The Internet's central mechanisms simply move information indiscriminately; the core of the network does not need to know if it is transmitting bits of data from an e-mail, a website, streaming audio, or some as-of-yet uninvented service.

The key to understanding the flexibility of the Internet is to realize that at the macro-level, the Internet can be thought of as constituting a *composite* standard,

9. In discussing simple, micro-level standards in the case of the Internet, I use the words 'protocol' (the common computer networking term) and 'standard' (the more general term) interchangeably.

with a whole series of simple standards as its component parts. The value of connecting to the Internet is not simply derived from coordination around TCP/IP as a way of exchanging data traffic. When people talk about the value of being on the Internet, they clearly have in mind something other than the micro-level mechanics of how their computers communicate with one another. Rather, network effects in the case of the Internet are derived from coordination around the entire package – standards that allow the exchange of e-mail, web browsing, streaming audio, encryption, and many more. The hypertext transfer protocol (HTTP) governing connection to the World Wide Web, for instance, was not a part of the Internet at its origins, but it is clearly an essential part of the Internet's composite-standard today. At the macro-level, therefore, the combination of parts that makes up the Internet's composite-standard can change significantly over time.

A simple model can clarify this point. Consider first, as a baseline for comparison, a network of 5 users of the QWERTY keyboard (Figure 1). Each of the typists gets value from the ability to sit down at another's machine and type just as they would on their own. We can calculate the value of the network by counting the number of direct connections between users: in this case, 10. The effect of increasing returns in this example is obvious when one considers increasing the number of users: six users would generate a value of 15, seven users would raise the value to 21, and in general, N users would generate a network with value [N*(N–1)]/2. Given these increasing returns, future typists will face an incentive to train on the QWERTY keyboard rather than an alternative keyboard configuration that has fewer adherents.

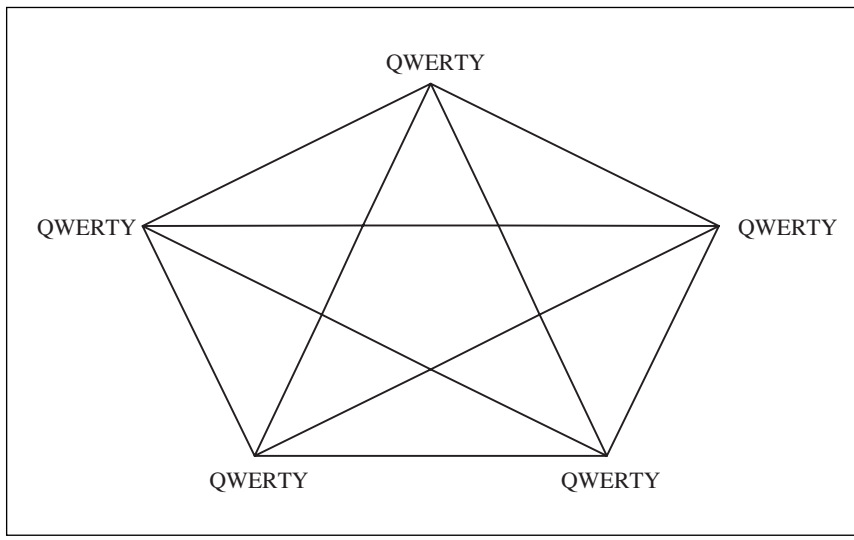

Figure 1. A Network with Simple Standards: QWERTY

Now consider how the nature of increasing returns changes when we model just a small portion of the complexity of the modern-day Internet, representing it not as a single, simple standard for exchanging data but as a composite of multiple simple standards that allow computers and their users to communicate in different ways. Figure 2 shows a model of the Internet in which the connection of computers via the core protocol TCP/IP is augmented by the presence of additional standards, allowing for web browsing and the exchange of email. These new standards still function through TCP/IP – hence the dotted lines running between them within each box – but they provide additional ways for computers to exchange data, and thus additional value for users of the network. If we weigh equally each of these new links between network users, the value of the network has increased to 30. Moreover, returns increase at a faster rate: with 3 standards, the value of the network would be [3*N*(N–1)]/2.

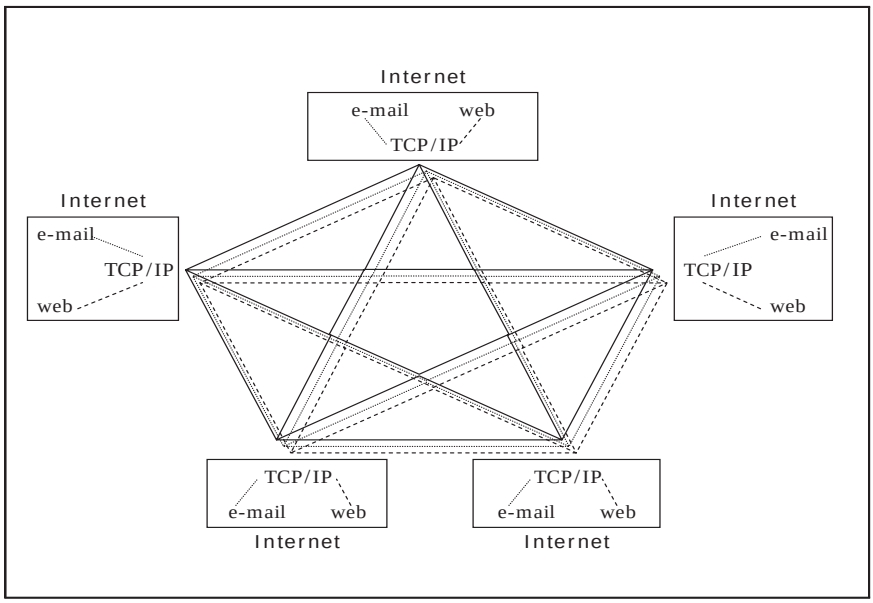

Figure 2. A Network with Composite Standards: The Internet

Despite the obvious quantitative change in the nature of increasing returns when we consider the multiple standards of the modern-day Internet, the crucial distinction between these two models is qualitative. While one can view the multiple parallel lines in Figure 2 as simply being multiple connections between users coordinating around three different standards, a more conceptually appealing way to interpret this

model is to consider these parallel lines to be thick single lines connecting users of a single composite-standard: the Internet. With change over time in the number and particular combination of constituent parts within this composite-standard, the thickness of the lines – and hence, the value of the Internet as a network – changes as well. More formally, one could use the parameter T (for 'thickness quotient') to represent the number of simple standards the network users have in common, and thus calculate the value of the network as [T*N*(N–1)]/2. The crucial insight, however, is that the fundamental nature of this network of composite-standards – what it does and does not allow users to do – can evolve significantly along with this changing collection of constituent parts.

The Internet's ability to evolve through the addition of new micro-level components has not only allowed it to become more valuable to users through the addition of services such as e-mail and the Web, but also to become more amenable to corporate and governmental control. The managers of corporate computer networks, for instance, often add features that allow them to monitor employees' usage of the Internet, even blocking access to non-work-related applications such as online music. The owners of broadband cable networks, which typically enjoy a local monopoly in their area of service provision, have been able to alter their networks to improve the speed and quality of Internet content from affiliated companies, so that users will be discouraged from visiting the web sites of their competitors (Bar et al., 2000). The governments of authoritarian regimes such as China and Saudi Arabia have used their control of national networks to block offending websites and monitor Internet traffic within their borders (Kalathil and Boas, 2003; Boas, 2006).

To return to our composite-standard model, let us consider the addition of a sixth network node under the control of an authoritarian regime, as depicted in Figure 3.[10] The ruler of this country faces an obvious incentive to connect to the existing Internet, which has a higher number of total users – and thus a higher overall value – than any separate computer network that might be built from scratch. But this ruler's interests are not served by the existing Internet, which, among other things, allows users to visit any website they want. Hence, the ruler adds additional control-facilitating features to the Internet for users connecting from within this country, specifically, a firewall that blocks certain websites. Unlike the addition of services such as e-mail and the Web, which permit all Internet users to exchange information in new ways, this national firewall does *not* increase the overall value of the Internet. Nonetheless, it does fundamentally alter the experience of using the Internet for citizens of this country, whose only option for using the Internet is to connect via a set of standards that allows for important elements of control. And as more and more control-seeking entities join the Internet and add their own control-facilitating features, the initially control-frustrating nature of this path dependent technology can be fundamentally transformed.

10. The multiple parallel lines in the previous diagrams are now replaced with single thick ones.

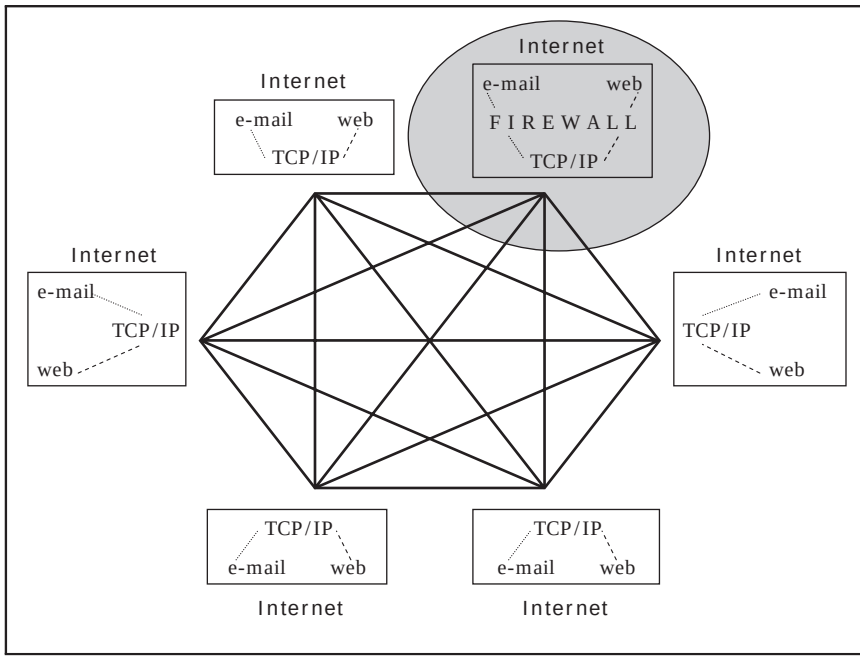

Figure 3. Implementing Control over a Portion of the Internet

If increasing returns at the macro-level of the Internet's composite-standard illustrate why some governments and corporations have sought to transform the existing Internet, increasing returns at the micro-level explain why this process of change has taken the particular form that it has. It is essential to note that the Internet's increasing amenability to control has come not through alteration of its underlying core protocols, but rather by adding new features that piggyback on top of these protocols, altering the nature of the Internet as a whole without replacing or reforming any of its existing parts. Simple standards such as TCP/IP are subject to increasing returns just like the composite of these standards; as such, wholesale replacement entails significant cost. But unlike the Internet at a macro-level, the Internet's micro-level standards cannot evolve through a changing mix of lower-level parts; they are relatively rigid and resistant to change.[11] Thus, because of the increasing returns at the level of inflexible, simple standards, actors seek to change the Internet by adding new control-facilitating features on top of those that already exist.

11. Changing these micro-level standards is not impossible. Indeed, TCP/IP has been upgraded several times already, and another revision is currently underway, in order to dramatically increase the number of computers with distinct network addresses that can be connected to the Internet. But the progress of this upgrade has been slow and uneven (Hovav and Schuff, 2005), and the nature of this change is fundamentally different from the flexibility that can be obtained by adding new services to the Internet.

In their relative resistance to change over time, the Internet's micro-level standards such as TCP/IP are somewhat akin to the simple standard of the QWERTY typewriter keyboard. The crucial difference between these two examples concerns the fact that only in the case of the Internet are these micro-level standards part of a larger composite-standard – the Internet at a macro-level – which itself generates increasing returns. This distinction is important because it explains why various improvements in the technology of typing devices, which in a certain sense have 'piggybacked' on top of QWERTY, have not generated a type of path dependence analogous to that of the Internet. Typing on a QWERTY keyboard has become markedly more efficient since its introduction, thanks to such innovations as electric typewriters, correction tape, and the personal computer (Williamson, 1996; Schwartz, 2004). But none of these efficiency related innovations constitutes additional standards around which typists coordinate. The owner of a QWERTY manual typewriter should not care if other typists are using correction tape or personal computers; the only element of a typing device that engenders coordination is the layout of the keyboard.[12]

In sum, the composite-standard model of path dependence illustrates how a path dependent process built around increasing returns can also accumulate fundamental change over time, through changes in the mix of constituent parts that generate these increasing returns. While the discussion above applied this model to a case of complex technological development, it can just as easily be used to explain the development of complex political institutions that engender coordination among political actors and are capable of evolution through a changing mix of constituent parts. In applying the composite-standard model to the political realm, we stand to gain a number of important methodological and theoretical insights into processes of institutional continuity and change.

The Composite-Standard Model and the Analysis of Political Institutions

As with the notion of a technological standard, the concept of 'institution' in political science is one that may involve multiple nested layers with varying degrees of complexity. For many rational choice scholars, the institutions incorporated into formal models are simple – micro-level institutions – agenda control, for instance – that prevents cycling and allows the formation of stable aggregate preferences.[13] Because they cannot be subdivided – agenda control is simply a rule, which must be either present or absent – institutions such as these are akin

12. Other capabilities of some of these typing devices, such as the ability to exchange files with compatible word processors, do engender additional coordination effects, but these effects have nothing to do with the keyboard used to enter text.

13. To be sure, there are also prominent examples of rational-choice scholars who examine broader historical processes and focus on macro-institutions (e.g. Bates et al., 1998; Acemoglu and Robinson, 2005).

to simple standards such as the QWERTY keyboard. By contrast, historical institutionalists typically deal with institutions at a higher level of aggregation: institutions are *composed* of individual rules and procedures, but the focus of inquiry is typically on the sum of these parts. The agenda control that might be incorporated into a spatial voting model, for instance, is a component of a legislative committee; such committees are in turn components of Congress, which is itself a part of the overall institution of government. Structurally, therefore, political institutions at a macro-level are more akin to composite technological standards, with a number of simple standards as their constituent parts.

Because the institutions of interest to historically oriented scholars are typically at the top of these nested hierarchies, the use of a QWERTY-inspired model of path dependence in historical institutionalism is something of a logical mismatch. Indeed, much of the recent research on institutional continuity and change has implicitly acknowledged this fact and has sought to distance itself from the notions of static institutional reproduction that is characteristic of QWERTY. Initial methodological statements on path dependence in political science and sociology (Mahoney, 2000; Pierson, 2000) built explicitly on the models of path dependence developed by economic historians and argued that the institutional stability of a path dependent process derived from increasing returns and positive feedback. More recent work on institutional development, however, has reacted to this initial emphasis on stability by arguing that political institutions often display greater change over time than is captured in existing path dependence models, and by introducing new mechanisms, such as layering and conversion, to account for this change.

In this section I argue that the specific contribution of the composite-standard model of path dependence to debates about continuity and change in political institutions is that it unifies such mechanisms as layering, conversion, and increasing returns within a single theoretical framework and highlights new ways in which these processes are intimately interconnected. In doing so, it builds upon and extends recent work by scholars such as Mahoney (2000), Pierson (2000), Schickler (2001), Schwartz (2004), and Thelen (2003, 2004), clarifying certain points of ambiguity in their discussions of institutional development. While recent research has tended to suggest that increasing returns explain only elements of institutional stability, for instance, the composite-standard model highlights the fact that increasing returns may actually be a necessary component of institutional change via layering or conversion. Moreover, the composite-standard model of path dependence helps to illustrate how both of these mechanisms, which have previously been posited as separate and distinct, can constitute two interrelated parts of a single case of institutional change.

The Path Dependence Debate: From Increasing Returns to Layering and Conversion

In an important statement in the recent debate on path dependence, Paul Pierson (2000, 2004) argues that increasing returns and continuity over time are as likely

to occur in the political realm as in the economy. Pierson maintains that the four economic characteristics generating increasing returns (fixed costs, learning effects, network effects, and adaptive expectations) are also common in political institutions, and he notes several political features that give rise to stability, such as the deliberate obstacles to change that may be inserted into national constitutions. Moreover, Pierson (2000: 263) sees a connection between increasing returns and stability, arguing that an increasing returns process 'may lead to a single equilibrium' which 'will in turn be resistant to change'. While later in the article he acknowledges the capacity for 'bounded change' within institutions (p. 265), one gets the sense from his argument that resistance to fundamental change is a *result* of the positive feedback driving a path dependent process.

James Mahoney's (2000) treatment of path dependence in historical sociology similarly maintains that the positive feedback mechanisms sustaining a self-reinforcing sequence often lead to institutional persistence. Mahoney's discussion goes beyond the utilitarian calculation of self-interest implied by increasing returns in economics and focuses on several additional mechanisms of institutional reproduction, including functional (where an institution is reproduced because of its function for an overall system), power (in which an institution empowers and is protected by an elite group), and legitimation (where an institution is reproduced because it is considered morally appropriate). What these various mechanisms of reproduction have in common is that they 'may be so causally efficacious that they "lock-in" a given institutional pattern, making it extremely difficult to abolish' (p. 515).

In his critique of the models of institutional path dependence elaborated by Pierson and Mahoney, Herman Schwartz (2004: 4–6) argues even more forcefully that 'formalized path dependence' implies a logical interconnection between increasing returns and a notion of stasis derived from punctuated equilibrium models, and he suggests that the concept of path dependence collapses when either of these components is removed. Schwartz maintains that increasing returns preserve the status quo, implying stasis; path dependence models must therefore relegate institutional change to critical junctures, rather than allowing for change within an increasing returns process. To support this claim, he cites Pierson's argument as summarized earlier and also notes that these characteristics (along with contingency) were all present in Arthur's (1989) formulation of path dependence models in economics.

The idea that increasing returns imply stability also underlies the recent formulation of several mechanisms that Thelen (2003) has proposed to account for the accumulation of change over time in political institutions. The first of these mechanisms is conversion, the process in which an existing institution is reoriented to serve new purposes or reflect new power dynamics (pp. 331–4). Such reorientation occurs either when a shift in external environment demands that an existing institution change in order to survive, or when the institution incorporates or is captured by a new constituency. Thelen points to an instance of conversion in the

development of the German vocational training system, which was originally designed to benefit artisans and pre-empt the influence of organized labor, but was fundamentally transformed when unions gained greater power and were incorporated into the system.

Layering is a second, distinct mechanism that has been proposed to account for change over time in the development of political institutions (Schickler, 2001; Thelen, 2003). In a layering process, an institution is changed incrementally as additional rules or structures are added on top of what already exists. Each additional layer constitutes only a small change of the institution as a whole, but this process can cumulate into an eventual transformation of the institution's fundamental nature. Typically, institutional evolution through layering results from actors' efforts to shape the development of an institution when they are either incapable of changing it at a macro-level or do not seek to do so. Schickler (2001), for instance, argues that incremental change in the US Congress results from competition among groups that do not have the strength to alter the underlying structures of Congress in ways that serve their interests. Rather, they add layers to the rules that already exist, implementing ad hoc solutions to immediate problems.

Like other contributors to the debate on path dependence, Thelen (2004) suggests that increasing returns imply resistance to fundamental change, and the mechanisms she proposes to account for such change are intended to take a clear step beyond the notion of increasing returns. She argues, for instance, that 'increasing returns arguments tell part of the story' of the transformation of German vocational training, 'but they are mostly designed to capture the logic of institutional reproduction, not institutional change' (p. 293). In other words, while Thelen certainly does not imply that institutional change is impossible when increasing returns are present, she does treat increasing returns as a mechanism that primarily explains the elements of continuity within a larger process of institutional development. Elements of change within this larger process are intended to be captured by the mechanisms of layering and conversion, which she characterizes as tools that are *distinct* from path dependence models based on increasing returns (Thelen, 2003: 231).

In sum, four of the scholars who have taken what are in other respects differing positions in recent debates on path dependence have all either explicitly or implicitly endorsed the notion that increasing returns principally explain elements of continuity in the development of political institutions, and those who have sought to explain elements of change have introduced new mechanisms such as layering and conversion, which are presented as being distinct from increasing returns. This conclusion about the explanatory limitation of increasing returns arguments is not surprising when one considers the role of increasing returns in models of path dependence derived from QWERTY. When increasing returns involve the reproduction of a simple, inflexible standard, there is no room for incremental change through mechanisms such as layering or conversion. If one is

limited to QWERTY-based notions of increasing returns, mechanisms for explaining institutional change would indeed have to be introduced as a separate set of tools that fall outside the realm of path dependence.

I argue, however, that by adopting the composite-standard model of path dependence to explain the development of macro-level political institutions, we can unify such mechanisms as increasing returns, layering, and conversion under a single theoretical roof and also highlight ways in which these mechanisms are intimately interconnected. Ultimately, increasing returns, layering, and conversion are not simply separate sets of tools that are united merely by virtue of being stored within the analyst's toolbox. On the contrary, they are more akin to interconnected cogs in a machine of institutional evolution, working together in a complex fashion to constitute real-world processes of continuity and change. If we could somehow remove one of these cogs, we would not simply be telling a smaller part of the story. Rather, the entire machine of institutional development might cease to function. Without increasing returns, the pursuit of layering and conversion might not make any sense. Moreover, the conversion of an institution may sometimes only be possible by means of a layering process. An application of the composite-standard model of path dependence can clarify these points.

The Composite-Standard Model: Unifying Layering, Conversion, and Increasing Returns

Unlike simple QWERTY-inspired models of path dependence, the composite-standard model can easily incorporate the mechanisms of layering and conversion when applied to a case of institutional change. In the application below, I return to the basic layout of the composite-standard model introduced in the previous section, substituting a generic political institution and its component procedures for the macro- and micro-level standards associated with the specific example of the Internet. As in the previous application of this model, I calculate the value of an institution (previously the value of the network) as the number of individual connections it allows between players: specifically, [T*N*(N−1)]/2, where N is the number of players, and T is a thickness quotient based on the number of micro-level standards (equally weighted) that engender coordination. In this case, however, I allow T to vary among players based on whether they actually obtain benefits from coordinating around a particular micro-level procedure. I treat the value of the institution for each player as the payoff that he or she obtains, and I assume that players will prefer institutional arrangements with higher payoffs. Finally, I assume that any player can change the existing macro-institution by removing or incorporating procedures, and that all players will then decide whether to continue coordinating around the revised institution, based on the change in institutional value that results from revision.

Consider first a situation in which Players 1, 2, 3, and 4 are committed to a macro-level Institution *I*, with Procedures *a*, *b*, and *c* as its constituent parts – specific

rules and procedures that each of the players has agreed upon. In this case, the value of the institution is 18. Now consider two additional Players, 5 and 6, who also benefit from coordinating around a macro-institution with Procedure *a* as one of its components, but only if Procedures *b* and *c* are removed, or, alternatively, an additional Procedure *d* is added (Procedure *d* could be thought of as a rule that counters the negative effects of *b* and *c* for Players 5 and 6). Like the element of control added to the Internet by an authoritarian leader in the previous example, Procedure *d* does not generate network effects, but if added to Institution *I* it would become a constituent part of the institution, just like the other three procedures. Moreover, if Procedures *b* and *c* are both components of Institution *I*, they generate no positive network effects for Players 5 and 6. Finally, in this case let *d* be a rule that benefits these two new players at the expense of the existing four, reducing the value of the institution by a fixed amount *r* for Players 1 through 4.

While the ultimate outcome depends on the value of *r*, it is quite likely that this game will result in a change in Institution *I* through the incorporation of the additional Procedure *d*. If Players 5 and 6 were to coordinate among themselves around an alternative institution *A* that incorporated Procedure *a*, the value of this alternate institution would be merely 1. However, Players 5 and 6 could also consider coordinating around the existing institution *I* if they were able to change it for their benefit, either by adding Procedure *d* (creating Institution I_1) or by eliminating Procedures *b* and *c* (creating Institution I_2). If either of these scenarios occurred, it would constitute an instance of conversion: new players would have captured an existing institution and reoriented it to serve their purposes. Moreover, if the ultimate outcome is one in which Procedure *d* is incorporated into Institution *I*, it will constitute an instance of layering, a new procedure being added to an institution and slightly transforming it as a result.

If Players 5 and 6 add Procedure *d* to Institution *I*, and the four original players remain committed to the institution, then the value of I_1 for Players 5 and 6 is 15, calculated on the basis of 6 players and the one standard a that generates benefits for the two newcomers. For Players 1, 2, 3, and 4, the value of the institution – based on six players coordinating around procedures *a*, *b*, and *c* – would be 45, minus the fixed amount *r* that is associated with the presence of Procedure *d*. If, on the other hand, Players 5 and 6 remove Procedures *b* and *c* from Institution *I*, all players are coordinating around an Institution I_2 that consists only of a single Procedure *a*, generating a value of 15 for every player. While Players 5 and 6 are indifferent between Institutions I_1 and I_2 but pre*f*er either of these to institution *A*, players 1 through 4 would consider Institution I_2(with a value of 15) to be inferior to the original Institution *I* (with a value of 18). Whether Players 1 through 4 would prefer I_1 over *I* would depend on the *r*, since the value of I_1 for these players is 45-*r*. However, there is clearly a large range of values of this parameter for which everyone's preferred outcome would be for Players 5 and 6 to reorient Institution *I* by incorporating Procedure *d*.

This generic application of the composite-standard model illustrates several key points about institutional change and continuity. First, increasing returns-based utility calculations are not only an important part of this model; they are crucial to explaining why conversion and layering happen at all. In other words, increasing returns in this particular application of the composite-standard model do not explain institutional stability; rather, they explain institutional change. In the absence of increasing returns-based utility calculations, the two newcomers would not have an incentive to pursue the conversion of Institution *I*, nor would the four existing players have any incentive to accept an alteration of the set of procedures that comprise this institution. Moreover, increasing-returns based calculations explain why the new players entering the game will try to achieve their objectives through a layering process – the addition of a new procedure rather than the removal of existing procedures. Current players would oppose efforts to alter Institution *I* by eliminating existing components from which they gain value. The only viable course of action for the newly interested parties is to add a layer to the existing institution in a way that serves their purposes.

Indeed, many real-life examples of path dependence in politics suggest that increasing returns-based calculations are crucial for explaining why institutional change via conversion and layering takes place. Precisely because of the increasing returns-driven value of name recognition and existing relationships, for instance, newly empowered political groups often seek to gain control of and reorient existing political parties rather than found new ones. This dynamic has led to fundamental change over time in a number of parties, from the Democratic Party in the United States to Mexico's Institutional Revolutionary Party, which were captured and converted by particular factions. Thelen (2003, 2004) acknowledges the same sort of cost-benefit calculation in the German labor movement's decision to reorient the existing vocational training system rather than destroy it.

In a similar fashion, increasing returns at the micro-level are often essential for explaining institutional change via layering. In Schickler's (2001) analysis of continuity and change in the US Congress, for instance, actors' efforts at layering are specifically a response to the path dependent replication of existing, inflexible rules. Only in rare instances like the Republican revolution of 1995 were reformers able to modify existing Congressional procedures that were protected by important constituencies. More commonly, groups that desired change found it less costly to layer additional procedures on top of existing ones.

Another key lesson that can be learned from this application of the composite-standard model of path dependence is that both layering and conversion, previously posited as separate and distinct mechanisms, may actually coexist in a single instance of institutional change. In the model presented earlier, layering is the means by which conversion ultimately occurs; the two processes are intimately interconnected. The compatibility of these two mechanisms is not surprising, since they occur at two different levels in the nested hierarchy of institutional standards. The decision to pursue conversion is based on a calculation of

the increasing returns-driven value of an entire institution at the macro-level, while layering efforts respond to the lock-in of components of this macro-institution. Thus, just as a political party, for instance, is composed of multiple rules and procedures, it also makes sense that the reorientation of this party via conversion occurs through the layering of addition rules and procedures that incrementally change its meaning and purpose.

An example of political change that illustrates both of these mechanisms simultaneously is the transformation of Mexico's Party of the Mexican Revolution into the Institutional Revolutionary Party in 1946 (Collier and Collier, 1991). At the macro-level, this process is a clear instance of conversion: a new faction gained control of the radical populist party crafted by Lázaro Cárdenas and reoriented it in a more conservative direction, creating the hegemonic centrist party that governed Mexico through the end of the 20th century. The process by which these changes were accomplished, however, contained important elements of layering. In particular, reformers tried to eliminate a system of candidate nomination by social sectors (such as labor and peasants), but the system was later reinstated after strong opposition from organized labor, an important constituency which benefited from this procedure and sought to protect it. Ultimately, the reorientation of the party was accomplished by creating a new layer of centralized hierarchical control and subordinating the old sectoral institutions beneath it rather than by eliminating these institutions entirely.

Indeed, the potential coexistence of layering and conversion is an idea that is consistent with Thelen's (2003, 2004) most recent work. While she first introduced the transformation of German vocational training as an exemplar of conversion, in her more extended treatment of this case she suggests that layering played an important role as well. Labor unions did not capture the German training system in one single dramatic episode. Rather, they established their own training system, which ran parallel for a time to the system initially created for artisans, but which also helped them eventually gain control of German vocational training as a whole. The specific advantage of applying the composite-standard model to this case is that it would highlight the logical interconnection between what Thelen simply presents as two separate mechanisms of institutional change.

Conclusion

Ongoing methodological discussions about the development of political institutions, with contributions from such scholars as Mahoney (2000), Pierson (2000), Alexander (2001), Thelen (2003, 2004), Schwartz (2004), Hacker (2004), and Crouch and Farrell (2004), constitute a productive intellectual exchange that has helped scholars refine arguments about historical causation in a systematic and rigorous manner. Recent scholarship is undoubtedly on the right track in shifting the emphasis from the ways in which institutions may be stable over time to the

manner in which they can change substantially through incremental evolution. In focusing an analytic lens on the mechanisms underlying institutional change, however, most of the contributors to this literature have sought to move away from increasing returns-based models of path dependence.

In this article I have offered a 'friendly amendment' to the arguments put forth by Mahoney, Pierson, Schickler, Schwartz, Thelen, and others, presenting the composite-standard model of path dependence as a tool that unifies the mechanisms of increasing returns, layering, and conversion, and shows how they can be intimately interconnected in instances of institutional change. It should be clear that I have no fundamental disagreements with their conclusions; indeed, I have used examples from Thelen's and Schickler's work to illustrate some of my revised interpretations of the ways that increasing returns, layering, and conversion work together. Rather than a refutation of existing arguments, I consider the composite-standard model to offer a useful extension of ideas presented in the ongoing discussion of institutional change and continuity.

Ultimately, the composite-standard model is a more general and potentially more accurate model of path dependence than those inspired by QWERTY, at the expense of the parsimony that the simple QWERTY-based model provides (Przeworski and Teune, 1970).[14] The composite-standard model is more general because it subsumes the logic of QWERTY: it combines the micro-level simple standards of which QWERTY is a case with the flexibility entailed by coordination around the sum of these component parts. For analyzing the evolution of complex technologies like the Internet, or complex processes of institutional change such as those examined by Thelen and Schickler, the composite-standard model of path dependence is also a more accurate model, capable of explaining elements of both continuity and change via increasing returns, layering, and conversion.

Methodological choices ultimately involve trade-offs (Collier et al., 2004), and the choice to adopt a composite-standard model of path dependence is no exception. In this case, the price of accuracy and generality is abandoning a simple and familiar heuristic device in favor of a more complex model, in which the interplay of change and continuity is perhaps more difficult to capture at first glance. Nonetheless, historical institutionalism and comparative historical analysis are traditions of inquiry that have typically favored complex but more complete explanations over those that are parsimonious but partial (Ragin, 1987; Pierson and Skocpol, 2002; Mahoney and Rueschemeyer, 2003). For those with an *a priori* commitment to tackling complex processes of institutional development and explaining both change and continuity over time, the composite- standard model of path dependence should prove appealing.

14. Przeworski and Teune (1970) suggest that more accurate theories are usually less general and less parsimonious. In this case, however, the composite-standard model represents a potential improvement over QWERTY-inspired formulations in terms of both generality and accuracy.

REFERENCES

Acemoglu, Daron and James A. Robinson (2005) *Economic Origins of Dictatorship and Democracy*. New York: Cambridge University Press.

Alexander, Gerard (2001) 'Institutions, Path Dependence, and Democratic Consolidation', *Journal of Theoretical Politics* 13(3): 249–70.

Arthur, W. Brian (1988) 'Self-reinforcing Mechanisms in Economics', in Philip W. Anderson, Kenneth J. Arrow and David Pines (eds) *The Economy as an Evolving, Complex System*, pp. 9–32. Reading, MA: Addison-Wesley.

Arthur, W. Brian (1989) 'Competing Technologies, Increasing Returns, and Lock-in by Historical Events', *Economic Journal* 99: 116–31.

Arthur, W. Brian (1994) *Increasing Returns and Path Dependence in the Economy*. Ann Arbor, MI: The University of Michigan Press.

Arthur, W. Brian, Yuri M. Ermoliev and Yuri M. Kaniovski (1983) 'On Generalized Urn Schemes of the Polya Kind', *Cybernetics* 19: 61–71.

Bates, Robert H., Avner Greif, Margaret Levi, Jean-Laurent Rosenthal and Barry R. Weingast (1998) *Analytic Narratives*. Princeton, NJ: Princeton University Press.

Bar, François, Stephen Cohen, Peter Cowhey, Brad DeLong, Michael Kleemen and John Zysman (2000) 'Access and Innovation Policy for the Third-generation Internet', *Telecommunications Policy* 24: 489–518.

Boas, Taylor C. (2006) 'Weaving the Authoritarian Web: The Control of Internet Use in Nondemocratic Regimes', in John Zysman and Abraham Newman (eds) *How Revolutionary Was the Digital Revolution? National Responses, Market Transitions, and Global Technology*, pp. 373–90. Stanford, CA: Stanford University Press.

Collier, David, Henry E. Brady and Jason Seawright (2004) 'Critiques, Responses, and Trade-offs: Drawing Together the Debate', in Henry E. Brady and David Collier (eds) *Rethinking Social Inquiry: Diverse Tools, Shared Standards*, pp. 195–228. Lanham, MD: Rowman & Littlefield.

Collier, Ruth Berins and David Collier (1991) *Shaping the Political Arena: Critical Junctures, the Labor Movement, and Regime Dynamics in Latin America*. Princeton, NJ: Princeton University Press.

Crouch, Colin and Henry Farrell (2004) 'Breaking the Path of Institutional Development? Alternatives to the New Determinism', *Rationality and Society* 16(1): 5–43.

David, Paul A. (1985) 'Clio and the Economics of QWERTY', *American Economic Review* 75(2): 332–7.

David, Paul A. (1986) 'Understanding the Economics of QWERTY: The Necessity of History', in William N. Parker (ed.) *Economic History and the Modern Economist*, pp. 30–49. New York: Basil Blackwell.

David, Paul A. (1987) 'Some New Standards for the Economics of Standardization in the Information Age', in Partha Dasgupta and Paul Stoneman (eds) *Economic Policy and Technological Performance*, pp. 206–39. New York: Cambridge University Press.

David, Paul A. (1999) 'At Last, a Remedy for Chronic QWERTY-skepticism!' Paper presented at the European Summer School in Industrial Dynamics, l'Institute d'Etudes Scientifique de Cargèse (Corse), France, 5–12 September.

David, Paul A. (2001) 'The Evolving Accidental Information Super Highway', *Oxford Review of Economic Policy* 17(2): 159–87.

Gerschenkron, Alexander (1962) *Economic Backwardness in Historical Perspective*. Cambridge, MA: Harvard University Press.

Hacker, Jacob (2004) 'Privatizing Risk without Privatizing the Welfare State: The Hidden Politics of Social Policy Retrenchment in the United States', *American Political Science Review* 98(2): 243–60.

Hovav, Anat and David Schuff (2005) 'The Changing Dynamic of the Internet: Early and Late Adopters of the IPv6 Standard', *Communications of the AIS* 15: 242–62.

Isenberg, David (1997) 'Rise of the Stupid Network', *Computer Telephony* (August): 16–26.

Kalathil, Shanthi and Taylor C. Boas (2003) *Open Networks, Closed Regimes: The Impact of the Internet on Authoritarian Rule*. Washington, DC: Carnegie Endowment for International Peace.

Lemley, Mark A. and David McGowan (1998) 'Legal Implications of Network Economic Effects', *California Law Review* 86(3): 481–511.
Liebowitz, S. J. and Stephen E. Margolis (1990) 'The Fable of the Keys', *Journal of Law and Economics* 33: 1–25.
Liebowitz, S. J. and Stephen E. Margolis (1995) 'Path Dependence, Lock-in, and History', *Journal of Law, Economics, and Organization* 11(1): 205–26.
Liebowitz, Stan J. and Stephen E. Margolis (2001) *Winners, Losers, and Microsoft: Competition and Antitrust in High Technology*. Oakland, CA: The Independent Institute.
Lipset, Seymour Martin and Stein Rokkan (eds) (1967) *Party Systems and Voter Alignments: Cross-national Perspectives*. New York: Free Press.
Mahoney, James (2000) 'Path Dependence in Historical Sociology', *Theory and Society* 29: 507–48.
Mahoney, James and Dietrich Rueschemeyer (eds) (2003) *Comparative Historical Analysis in the Social Sciences*. New York: Cambridge University Press.
Moore, Barrington, Jr. (1966) *Social Origins of Dictatorship and Democracy: Lord and Peasant in the Making of the Modern World*. Boston: Beacon Press.
Pierson, Paul (2000) 'Increasing Returns, Path Dependence, and the Study of Politics', *American Political Science Review* 94(2): 251–67.
Pierson, Paul (2004) *Politics in Time: History, Institutions, and Social Analysis*. Princeton, NJ: Princeton University Press.
Pierson, Paul and Theda Skocpol (2002) 'Historical Institutionalism in Contemporary Political Science', in Ira Katznelson and Helen V. Milner (eds) *Political Science: The State of the Discipline*, pp. 693–721. New York: W. W. Norton.
Przeworski, Adam and Henry Teune (1970) *The Logic of Comparative Social Inquiry*. New York: Wiley-Interscience.
Ragin, Charles C. (1987) *The Comparative Method: Moving Beyond Qualitative and Quantitative Strategies*. Berkeley, CA: University of California Press.
Saltzer, J. H., D. P. Reed and D. D. Clark (1984) 'End-to-end Arguments in System Design', *ACM Transactions in Computer Systems* 2(4): 277–88.
Schickler, Eric (2001) *Disjointed Pluralism: Institutional Innovation and the Development of the U.S. Congress*. Princeton, NJ: Princeton University Press.
Schwartz, Herman (2004) *Down the Wrong Path: Path Dependence, Increasing Returns, and Historical Institutionalism*. Unpublished Manuscript, University of Virginia.
Stinchcombe, Arthur (1968) *Constructing Social Theories*. New York: Harcourt, Brace, and World.
Thelen, Kathleen (2003) 'How Institutions Evolve: Insights from Comparative Historical Analysis', in James Mahoney and Dietrich Rueschemeyer (eds) *Comparative Historical Analysis in the Social Sciences*, pp. 208–40. New York: Cambridge University Press.
Thelen, Kathleen (2004) *How Institutions Evolve: The Political Economy of Skills in Germany, Britain, the United States, and Japan*. New York: Cambridge University Press.
Williamson, Oliver E (1996) *The Mechanisms of Governance*. New York: Oxford University Press.

TAYLOR C. BOAS is a PhD candidate in political science at the University of California, Berkeley. His dissertation analyzes the evolution of party–society linkages and political cleavages in Chile, Peru, and Brazil through a study of two decades' worth of presidential election campaigns in each country. His prior research examined the impact of the Internet on authoritarian rule. Recent publications have appeared in *Latin American Research Review* and *Studies in Comparative International Development.* ADDRESS: Department of Political Science, University of California, 210 Barrows, Berkeley, CA 94720–1950, USA [email: tboas@berkeley.edu].