Internet Architecture and Innovation

Barbara van Schewick

The MIT Press
Cambridge, Massachusetts
London, England

Set in Stone Sans and Stone Serif by Toppan Best-set Premedia Limited. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Van Schewick, Barbara.
Internet architecture and innovation / Barbara van Schewick.
p. cm.
Includes bibliographical references and index.
ISBN 978-0-262-01397-0 (hardcover : alk. paper) 1. Internet. 2. Computer network architectures. 3. Technological innovations. 4. Business—Data processing. I. Title.
TK5105.875.I57V378 2010
004.6'5—dc22

2009037130

10 9 8 7 6 5 4 3 2 1

Introduction

In the thirty years since its inception, the Internet has experienced remarkable growth. What started in the 1970s as an experimental network connecting research and military networks has become a global network linking more than 1.6 billion users worldwide.[1] The Internet's growth has been fueled by an unprecedented amount of innovation. Over the years, network engineers have developed numerous new physical networking technologies (including Wi-Fi and optical networking technologies) over which the Internet can run. The Internet now connects everything from sensors to supercomputers. A constant stream of new applications lets users do new things, or do them more efficiently, making the Internet more attractive and useful to diverse users. As more and more users adopt them, these applications, the content they help produce and make available, and the new economic, social, cultural, and political practices they enable are transforming all areas of society. They change, for example, how firms can organize themselves, make their products, and interact with customers and with other firms. They increase the opportunities available to us, helping us to be more productive in our professional and private lives; to interact with relatives, friends, and strangers; to get to know them, communicate, or work with them; to educate ourselves using a variety of sources; and to participate in social, cultural, and democratic discourse.

The Internet's growing size, its transition from a research network operated by public entities to a commercial network operated by commercial providers who need to make profits, and its transition from a network connecting a small community of users who trust one another to a global network with users who do not know one another and may even intend to harm one another put pressure on the Internet's technical foundations. To deal with the impending shortage of Internet addresses, users and network providers deploy network-address translators that let several devices share the same address. To protect their networks against attacks,

organizations put firewalls that block potentially harmful applications at the borders of their private networks. To increase their profits, network providers use technologies that enable them to identify and control the applications and the content passing through their networks.

These changes violate the design principles on which the Internet originally was based. They are hotly contested in technical forums such as the Internet Engineering Task Force, which develops and maintains the core Internet standards. The debate has also reached the political arena: arguing that these changes will destroy the Internet as a platform for innovation, free speech, and economic growth, scholars, public-interest organizations, industry organizations, and companies have asked regulatory agencies and legislative bodies all over the world to step in and restore the beneficial environment for innovation that the architecture of the Internet originally created.[2] In the United States, the debates over open access to broadband cable networks, over network neutrality, and over network management in broadband networks have all been shaped by such arguments.[3]

Independent of attempts to solve problems within the framework provided by the existing Internet, a new generation of networking research projects has begun to contemplate how the Internet of the future should look. Many of these projects start from a clean architectural slate, essentially creating a new Internet architecture from scratch.

In each of these cases, network engineers, legislators, and regulators face the same questions: How important are the design principles that shaped the Internet's original architecture? Are the (proposed) deviations much-needed improvements that reflect changed technical or business requirements, or are we losing the very features that were central to the Internet's success?

For regulators, the situation is even more complicated. Shouldn't a technical issue such as how to best structure the architecture of the Internet be left to network engineers? And if certain architectures foster innovation more than others, why won't the market create them?

This book aims to answer these questions. It examines how changes in the Internet's architecture (that is, its underlying technical structure) affect the economic environment for innovation and evaluates the impact of these changes from the perspective of public policy.

The answers are not obvious. After decades of research on innovation, we understand how changes in laws, in norms, or in prices affect the economic environment for innovation and how they affect innovators' decisions to innovate. We lack a similar understanding of how architecture affects innovation. Just as the architecture of a house describes its basic

inner structure, the architecture of a complex system describes the basic inner structure of the system—its components, what they do, and how they interact to provide the system's functionality. That such a technical structure may have economic consequences at all is a relatively recent insight.[4] Most people still think of architectures as technical artifacts that are relevant only to engineers.

Thus, understanding how the Internet's architecture affects innovation requires us to think more generally about how architectures affect innovation. How can the architecture of a complex system influence the economic system in general, or an economic activity such as innovation in particular? Which features of an architecture affect the economic environment for innovation, and how do these features vary across different architectures or design principles?

The answers to these questions are relevant beyond the Internet. Since in most cases a system's functional requirements determine the system's architecture only in part, system architects have considerable latitude to choose different architectures, and this latitude allows them to consider other goals too. This freedom is particularly pronounced in software-intensive systems that are less constrained by the laws of nature than the design of physical products is. In particular, if architectures have economic consequences, system architects can create architectures that prioritize and realize particular economic goals.

These insights open new opportunities for businesses, for law, and for public policy. Businesses may want to engage in "strategic design" by creating architectures that shape the competitive environment in their favor. In the future, being able to design architectures that further a firm's strategic interests or knowing how to evaluate other firms' architectural strategies and react to them may be as important to a firm's success in the marketplace as a firm's ability to engage in more conventional forms of competition.[5]

For law and public policy, the economic impact of architecture seems to be empowering and challenging at the same time. Traditionally, policy makers have used the law to bring about desired economic effects. Architecture may provide an alternative way of influencing economic systems. Apart from using architecture to realize their own economic goals, policy makers may have to constrain the extent to which private actors can use architecture to further their private economic interests. This is particularly relevant to communications policy, a field in which certain architectures may seriously restrict regulators' ability to regulate at a later stage. Moreover, as communications networks continue to permeate more and more

sectors of the economy, the negative effects of an architecture that strongly favors a few economic actors may be particularly long-term and severe.

To exploit the effect of architecture on the economic system in practice, to design architectures that further our interests, or to understand what other people's architectures may mean for us we need to understand exactly how an architecture influences the economic system and what features of an architecture we must tweak to create a specific economic effect. This book is a step toward that goal.

Overview of Chapters

Chapter 1 introduces a theoretical framework for understanding how architectures relate to economic systems and, more specifically, for understanding how architectures affect innovation. It explains architectures and design principles as concepts, and it highlights ways in which architectures influence economic systems and are influenced by them. In particular, the relationship between architectures and the economic system can be understood within a broader framework that economists use to explain the evolution of the economy as a whole. In this framework, the economic system evolves as economic actors pursue their own interests within a set of constraints, and as they act to change those constraints. Constraints delimit the options available to economic actors and influence the costs and benefits associated with these options. Well-known constraints include prices, laws, norms, and the natural and technical environments in which economic actors exist. Like these other constraints, an architecture can affect human behavior by imposing constraints on those who interact with the architecture or are exposed to it. Specifically, by imposing constraints on those who design, produce, and use a complex system, the architecture of the system (and the design principles that were used to create it) can influence the economic system in which the system is developed, produced, and used. Different architectures may impose different constraints, which may result in different decisions by economic actors, which in turn may result in different firm and market structures and different levels of economic activity. And by changing existing architectures or creating new ones, economic actors can change the constraints that architectures impose.

Before we can understand how specific architectures constrain economically, we need to understand how they constrain technically. Part II of the book introduces the architecture we will use to study the effect of architecture on innovation—the original architecture that governed the

Internet from its inception to the early 1990s[6]—and the design principles that were used to create that architecture: modularity, layering, and the end-to-end arguments.

While network engineers agree that the end-to-end arguments are among the few architectural principles underlying the architecture of the Internet,[7] other scholars have offered widely differing and often contradicting views on what the end-to-end arguments are, what they say, and how they relate to the architecture of the Internet. For example, some suggest that the end-to-end principle is "an important architectural principle that has governed the Internet since its inception"[8]; others say that "the end-to-end principle simply does not dictate a robustly specified functional design for the network"[9] but "follows from (and is an articulation of) the implicit design principle inherent to the layers model of the TCP/IP protocol [*sic*]"[10]; still others argue that "the Internet was never wholly end-to-end"[11] and that the end-to-end argument "is not an organizing principle; . . . if it is a principle, it is probably not true, and . . . even if it is true, it is probably not useful."[12] In policy debates concerning the architecture of the Internet, opponents of regulation often argue that proponents of certain regulatory interventions (for example, of open-access rules or network-neutrality rules) have stretched the end-to-end principle beyond its original meaning.[13] Even networking engineers often disagree about whether a certain technical solution violates the end-to-end arguments or not. In view of the high level of confusion, I discuss the end-to-end arguments and their relationship to the architecture of the Internet in detail.

My analysis yields an important insight: there is no single version of the end-to-end arguments, but two different ones that embody different rules for architectural design. The first version, which I call *the narrow version*, was presented by Jerome Saltzer, David Reed, and David Clark in the 1981 paper in which the end-to-end arguments were first named and identified as a design principle[14]; the second version, which I call *the broad version*, is the focus of later papers by these authors.[15] The difference between the two versions is not immediately apparent, and Saltzer, Reed, and Clark never explicitly drew attention to the change in definition. There are, however, real differences in scope, content, and validity that make it necessary to distinguish between the two versions. At the same time, the silent coexistence of two different design principles under the same name explains some of the confusion that surrounds the end-to-end arguments.

Chapter 2 describes the design principles that were used to create the original architecture of the Internet and highlights the trade-offs involved in each of them. Chapter 3 describes how these design principles shaped

the original architecture of the Internet. Chapter 3 also discusses some of the misconceptions surrounding the end-to-end arguments and their relationship to the architecture of the Internet. For example, scholars have argued that the end-to-end arguments rule out architectures that store state in the network, prohibit the provision of Quality of Service in the network,[16] require the network to be simple, or constrain the design of applications by preferring certain application-management structures over others. Others have argued that the original architecture of the Internet was not based on the end-to-end arguments on the ground that the end-to-end arguments are no different from the layering principle, or on the ground that the original paper describing the end-to-end arguments was published after the original architecture of the Internet was developed, or on the ground that e-mail had a certain management structure.

To evaluate how the Internet's architecture affects innovation, we must understand exactly how architecture constrains economic actors. How do design principles and the architectures they shape influence the costs and benefits associated with a given innovation? How do they affect who can design and build an innovation? Answering these questions is the goal of part III. Each chapter in that part focuses on a particular aspect of the relationship between architecture and innovation and explores it in detail. The resulting insights are then applied to assess how the Internet's original architecture affects innovation. Although both versions of the end-to-end arguments shaped the original architecture of the Internet, only the broad version affects the economic environment for innovation. Therefore, parts III and IV focus on modularity, layering, and the broad version.

Chapter 4 explores how architecture influences the cost of realizing an innovation. It first discusses the effect of architecture on the costs of innovation in the context of modular and integrated architectures. Using the theory of real options, it explores the relationship between the costs of innovation implied by different architectures and the value of experimentation under uncertainty. It shows how the different costs of change in modular and integrated architecture affect the option value of these architectures, and how these differences affect the willingness of component designers to take risks when designing their components.

Chapter 4 also explores how differences between alternative modular architectures influence the option value of comparable modules in these architectures. In particular, the analysis shows how differences in the costs of innovation with respect to a module, the nature of dependencies on that module, and the uncertainty surrounding a module may affect the level of investment and rate of change of that module in alternative archi-

tectures. The chapter then uses these general insights about the effect of architecture on the costs of innovation and on the likelihood of innovation to determine how specific design choices in the Internet's original architecture affect these factors. It explores how the decision for or against the broad version of the end-to-end arguments influences the costs of developing applications, and how the decision for or against the use of relaxed layering with a portability layer affects the costs of developing the physical network technologies over which the Internet can run. It also highlights how these architectural choices influence the costs of producing, distributing or deploying innovations.

The final subsection of chapter 4 explores how the Internet's original architecture affects the rate of change at the Internet and transport layers. In recent years, the difficulties associated with evolving the core of the Internet have received much attention in the networking community. My analysis traces these difficulties to the deployment context of the commercial Internet, which has made it much more difficult to change certain parts of the Internet's architecture than a purely technical analysis of the structural dependencies within the architecture would suggest. On a more abstract level, the analysis shows how the deployment requirements implied by a network architecture interact with the economic system in which the network is used to constrain the evolution of the architecture beyond what a real-options analysis would predict.

Chapter 5 focuses on the implications of architectural choices for the organizational structures in which the development and production of the system—and subsequent innovation—can take place. The analysis shows how an architecture affects how a single firm that has designed an architecture can organize the subsequent development and production of components to this architecture, and how the organizational options enabled by an architecture may ultimately influence the structure of the industries surrounding the architecture. Ultimately, these mechanisms determine whether independent economic actors other than the system's architect can develop and produce new components for the system, both initially and later. The chapter then applies these insights to the original architecture of the Internet. In particular, it discusses how that architecture affects who can develop new applications. In a network architecture based on the broad version of the end-to-end arguments, anyone who knows (or is willing to learn) how to program and has access to a computer connected to the Internet can develop new applications. An innovator does not have to be an employee of a firm or get outside funding to realize his or her idea for an application. As a result, the set of potential innovators is much

larger and more diverse than in network architectures that deviate from the broad version.

Chapter 6 explores how an architecture may affect the competitive strategies available to makers of complementary components. In particular, it analyzes how differences in a component's ability to monitor and control the execution of neighboring components may affect the strategies available to the maker or the users of the controlling component, and how this influences what benefits makers of complementary components can expect to realize. Although the chapter explores the effect of a component's ability to control other components in the context of the Internet's original architecture, the underlying theoretical framework is general enough to apply to other architectures. The chapter shows how architectural differences between the Internet's original architecture and a hypothetical architecture deviating from the broad version of the end-to-end arguments alter network providers' strategic options. It explores the conditions under which a network provider might have an incentive to discriminate against some applications, and it discusses recent instances of discrimination. This chapter also examines whether differences in market structure or laws may moderate the impact of architectural differences. It shows that network providers may have an incentive to exclude applications even if they face competition from other providers. The final section explores the effect of architectural differences on network providers' pricing strategies, such as their ability to charge users application-specific prices or to charge the providers of applications or content for the right to gain access to the customers of an Internet service provider.

Applications are the transmission belts that transform the general functionality of the Internet—transporting data packets from one computer to another—into something that creates value for users (and, in the process, creates value for society). Part IV of the book analyzes how changes in the architecture of the Internet affect the economic environment for innovation in applications and evaluates the public-policy consequences.

Chapters 4–6 provide a variety of insights into the relationship between architectures and design principles, on the one hand, and the economic system, on the other hand, both on a general level and with respect to the original architecture of the Internet and the design principles that shaped it. Chapter 7 draws on those insights to summarize how increases in the amount of application-specific functionality in the network's core change the economic environment for innovation at the application level. Thus, the chapter provides a framework in which the effects of deviations from

the broad version of the end-to-end arguments on application-level innovation can be assessed. It shows that as the amount of application-specific functionality in the network increases, independent application developers' incentives to innovate are reduced. At the same time, the range of potential innovators decreases, while the locus of control over the development and deployment of innovations in applications gradually shifts to the network owner. In addition, the chapter evaluates how such architectural changes affect other characteristics of the economic environment for application-level innovation, including network owners' incentives to innovate at this level and the costs associated with adopting new or improved applications.

Chapter 8 explores the effect of these differences on the amount of application innovation that will occur under different network architectures. In particular, drawing on theories of innovation and the histories of specific applications, chapter 8 examines how differences between decentralized and centralized environments for innovation and the associated differences in the number and types of innovators may affect the amount, the quality, and the character of application-level innovation. In a network architecture based on the broad version of the end-to-end arguments, anyone with access to an end host and with the ability to learn a programming language is a potential innovator. Innovators independently decide on their approaches to innovation; users independently choose which applications they want to use. In contrast, in architectures deviating from the broad version, network providers control which applications can be realized and used. At the same time, the set of potential innovators is smaller and less diverse; in the extreme case, network providers are the only remaining innovators. If there is uncertainty or consumer heterogeneity, a larger and more diverse group of innovators under an end-to-end architecture will produce more and better applications than a few network providers who control which applications are realized and used. If users' needs are heterogeneous, innovators in the end-to-end architecture will also produce more diverse applications that better meet users' needs.

Chapter 9 assesses how the various network architectures relate to the public interest and to network providers' private interests. Whereas earlier chapters focused on how network architectures affect innovation, this chapter looks more broadly at the economic and non-economic consequences of complying with or deviating from the broad version of the end-to-end arguments. The analysis shows that the public's interests in network architecture diverge from the network providers' interests,

creating a market failure regarding the evolution of the Internet's architecture.

The concluding chapter summarizes the book's main arguments and discusses the implications of the book's results for public policy and network design.

Overall, the book shows that the Internet's original architecture was based on a design principle that I call *the broad version of the end-to-end arguments*. This design fostered innovation in applications. Today, the Internet's architecture changes in ways that deviate from this principle. These changes reduce the amount and the quality of application innovation at significant costs to society: As a general-purpose technology, the Internet does not create value through its existence alone. It creates value by helping users do what they want to do, or by letting them do so more efficiently. Applications are the tools that let users realize this value. By reducing innovation in applications and limiting users' ability to decide how to use the network, these changes significantly reduce the Internet's usefulness and value for users, and, ultimately, for society as a whole. In addition, reducing innovation in applications limits the Internet's ability to contribute to economic growth. Finally, the Internet's ability to enhance individual freedom, its ability to provide a platform for better democratic participation, and its ability to foster a more critical and self-reflective culture are tightly linked to features resulting from the broad version of the end-to-end arguments. By removing these features, the changes to the architecture also threaten the Internet's ability to realize its social, political, and cultural potential. Though the broad version of the end-to-end arguments also has social costs, they are not large enough to justify sacrificing the social benefits created by the broad version.

While public interests suffer, network providers benefit from the changes. They control the evolution of the network, and it is highly unlikely that they will change course without government intervention.

Navigating the Book

The author of an academic book usually assumes that her readers know the relevant literature, and that she only has to add to it. I do not have this luxury. My argument crosses a number of disciplines and speaks to readers with a wide range of backgrounds. The book is designed to be accessible to all of them. In particular, one does not have to have a background in networking or in economics to understand it. As a result, the book is longer than it would be if it were targeted at readers in a single discipline, and it contains explanations that will be familiar to some

Table I.1
The end-to-end arguments from a technical perspective.

Understanding the End-to-End Arguments	Pages
What are the two versions of the end-to-end arguments?	57–75, 377–379
What are the differences between the two versions, and why should we distinguish between the two?	75–81
What is the trade-off underlying the broad version?	68–75, 355–371
Does the broad version differ from the layering principle?	104–105
Does the broad version prohibit quality of service?	106–107
Does the broad version require the network to be stupid or simple?	107
Do the end-to-end arguments prevent state in the network?	105–106
Does the broad version make it more difficult (or even impossible) to make the network secure?	366–368
How can you deviate from the broad version? What does it mean to make the network more controllable or more opaque?	286–287
The End-to-End Arguments and the Architecture of the Internet	**Pages**
How have the two versions shaped the Internet's original architecture?	90–103, 110–112, 379–381
Is the Internet Protocol general enough?	142, box 4.3
How do the two versions relate to current developments in the Internet, such as the evolution of applications toward a more distributed structure or the proliferation of middleboxes?	383–387, 371–372
Is the trade-off underlying the broad version still appropriate today?	368–371, 355–356
Should the end-to-end arguments continue to guide the Internet's architecture as a technical design principle?	388–389

readers but not to others. For example, engineers will already know about modularity, the layering principle, and the architecture of the Internet, and lawyers, economists, and management scientists will already know how transaction costs influence the boundary choices of firms. I hope that my headings and introductions will help readers recognize things they already know and help them get to what they do not know.

The book can be read straight through, from cover to cover. But one may also follow some of the more specific conceptual threads that run through the book. The first thread explores the end-to-end arguments from a technical perspective (table I.1). The second thread analyzes the social value of architectures based on the broad version of the end-to-end arguments and the resulting policy implications (table I.2).[17] The third thread explores aspects of the approach to "architecture and economics" (box I.1) advanced by the book. Within the framework described in

Table I.2

The social value of the broad version of the end-to-end arguments.

The End-to-End Arguments and Their Relationship to the Internet's Original Architecture	Pages
What are the two versions of the end-to-end arguments, and why should we distinguish between the two?	57–81, 377–379
How have they shaped the Internet's original architecture?	90–103, 110–112, 373–381
The Effect of the Broad Version on Application Innovation	**Pages**
How does the broad version affect the cost of developing new applications?	137–148, 383–387
How does the broad version affect who can develop new applications?	204–213
Can innovators with no or little outside funding really produce successful applications?	204–213, 300–308, 334–345
How does it affect independent application developers if network providers can control the applications on their network?	215–281
Do network providers have an incentive to discriminate against or exclude unaffiliated applications?	218–273
Does competition among network providers remove this incentive?	255–264
What pricing strategies can network providers follow in a network where they control the applications?	217–218, 273–280
How does the economic environment for application innovation change if networks deviate from the broad version?	285–295, 383–387
How do differences in the size and diversity of the innovator pool and differences in control over application innovation and deployment affect the overall amount, quality, and type of innovation?	297–351
How important is it if innovators other than the network provider can independently choose their approach to innovation?	345–348
How important is user choice for application innovation?	349–351
How do the architectural differences among network architectures that are based on the broad version or deviate from it, taken together, affect the overall amount, quality, and type of application innovation under these architectures?	351–353
The Social Value of Network Architectures Based on the Broad Version	**Pages**
What are the social benefits associated with the broad version?	355–365
How important is application innovation?	355–360
How important is user choice?	349–351, 361–363
How does the broad version affect the Internet's ability to realize its political, social and cultural potential?	364–365
What are the social costs associated with the broad version?	365–368
How should we trade off among the social benefits and social costs?	368–371
Do network providers' private interests diverge from the public interest?	371–375
Policy Implications	**Pages**
What do these insights mean for public policy?	218–221, 264, 272–273, 387–392
Should legislators mandate compliance with the broad version?	388

Box I.1
Architecture and Economics

This book's approach to the study of architecture and innovation is an example of a more general approach to studying the architecture of complex systems, an approach I call *architecture and economics*. The approach understands architecture as one of several constraints on human behavior and uses economic theory (broadly defined) to explore the effect of these constraints. In its narrowest meaning, "architecture and economics" denotes efforts to understand how the architectures of complex systems influence, and are influenced by, the economic systems in which the complex systems are designed, produced, and used. As we saw above, the links between architectures and economic systems have important implications for how businesses compete and how public policy is made. While this book focuses on the effect of architecture on a specific economic activity (innovation), the underlying framework is general and can be used to understand the effect of architecture on a much broader range of human behavior. Viewed from this perspective, "architecture and economics" describes a much broader field of research—efforts that use economic theory to understand how architectures affect specific forms of human behavior, and, more generally, how architectures influence, or are influenced by, economic, social, cultural, or political systems.[a] I use the term in the broader sense.[b]

a. I do not mean to imply that this type of research has not existed so far. Most research in this area, however, focuses on specific architectures or design principles. For example, a large body of literature in management science and engineering explores the economic effect of modular and integrated architectures in the design of physical products. (See note 4 to this chapter for references to this literature.) There is some research in software engineering that examines the economic effect of software architectures (e.g., Sullivan et al. 1999; Boehm and Sullivan 2000; Sullivan et al. 2001; Erdogmus et al. 2002). Representative examples of legal scholarship exploring the effect of the Internet's architecture are Lessig 1998, Lemley and Lessig 1999, Lessig 1999a, Lessig 2001, Wu 2003a, Benkler 2006, Balkin 2008, and Zittrain 2008. There is, however, no established field of research that connects work in this field under a common umbrella, nor is there an accepted framework or methodology for exploring these issues. Among others who have advocated for a more unified approach to the study of the architectures of complex systems are Baldwin and Clark (2006b) (who advocate a "science of design" that would cover questions very similar to the ones outlined in the text) and van Schewick (2004).
b. "Law and Economics" has a similar dual meaning. It is used to describe the study of the relationships between law and the economic system, but it also describes, more broadly, efforts to understand the effect of law on human behavior using economic theory. For an overview of the different schools of thought within this field, see Mercuro and Medema 2006.

Table I.3
Architecture and economics.

Factors Mediating the Effect of Architecture	Pages
How do non-architectural constraints interact with architecture?	
Theory	26–28, 31
Examples	196–197, 218–221, 264, 272–273
How do characteristics of the actors exposed to the architecture affect the impact of architecture?	
Theory	30
Examples	204–213, 298–345
How do actors' relationships with other actors affect the impact of architecture?	
Theory	31
Examples	133, 218–221, 255–264, 212–213
The Effect of the Economic System on Architecture	
How is architecture influenced by economic systems?	
Theory	3, 23–26, 28, 32, 389
Examples	151–163, 371–372, 389–392
The Effect of Architecture on Activities Other Than Innovation	
How can the architecture of a complex system affect the economic, social or political systems in which the complex system is used?	
Theory	28, 361–362
Examples	359–361, 362–365
Architecture and Economics: Implications	
What are the implications of the economic impact of architecture for businesses?	
Theory	3
Examples	371–372, 389–392
What are the implications of the economic impact of architecture for law and policy?	
Theory	3–4
Examples	388

chapter 1, the book focuses on the effect of one constraint (architecture) on one specific activity (innovation), and sets aside consideration of other factors (such as the effect of non-architectural constraints, or the mechanisms by which architectures are influenced by economic systems) that are relevant within the framework. This thread ties together the portions of the book that touch on these other factors (table I.3). In doing so, it complements the detailed analysis (in chapters 4–8) of how architectures—particularly that of the Internet—affect innovation.