14

Collaborative Architectures

Scott E. Page

Abstract

Collaborations serve a variety of ends and purposes. Most collaborations can be classified as one-shot or ongoing. The architectures within which collaborations occur (i.e., the rules, roles, processes, information structures, and incentives) matter. Having the right people with the requisite skills is not sufficient, particularly for complex tasks. Given the enormous set of possible architectures, finding an effective one requires forethought along with knowledge of the culture as well as the task at hand. Within a collaborative architecture, two types of adaptation occur: the participants learn and, often, the architecture adapts. In view of this, we should consider collaborations as coevolutionary processes.

Introduction

The word collaboration evokes images of small groups of people sitting in a room or around a table working together on a task. Perhaps, John, Paul, George, and Ringo proposing lyrics and melodies; Frida and Diego critiquing each other's paintings; or Daniel and Amos puzzling through deviations from rationality. Such images conjure up a situation where there are no fixed rules, just individual people sharing ideas and thoughts and offering improvements and refinements.

Processes of collaboration, though, encompass a far wider range of activities. Collaborations can involve hundreds if not thousands of participants. They can integrate ideas from participants in different locations who interact synchronously, such as groups of scientists who collaborate within and across organizations to produce vaccines and life-saving molecules. They can consist of people who belong to distinct firms in an automobile supply chain that collaborate to produce vehicles, or people who do not even know one another who collaborate to write Wikipedia articles and organize Reddit communities.

From "The Nature and Dynamics of Collaboration,"
edited by Paul F. M. J. Verschure et al. Strüngmann Forum Reports, vol. 33,
Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262548144

Collaborations occur between humans and other animals, humans and artificial intelligence, and humans and the built environment. A collaboration can occur entirely within the minds of participants, or the workings of processors, or it can be aided by artifacts. White boards, Slack channels, and other memory channels often play key roles in a collaboration.

All these types of collaborations, even the informal ones, rely on architectures: A *collaborative architecture* determines who participates and how their interactions are structured. It delineates roles, responsibilities, and communication protocols as well as the assignment of credit and enforcement of blame. It establishes rules for decision making and, often, meta rules or procedures for adapting the architecture.

This last attribute is particularly important, given the difficulties in designing an effective, much less an optimal or efficient architecture and the likelihood that the reason for the collaboration, and hence its architecture, may change over time. Consider, for a moment, child-rearing: parents, families, and friends encounter multiple challenges as they collaborate to raise a child. As the child ages, the challenges change as does the architecture needed to manage this new stage of the collaboration. Related phenomena occur within most collaborations. As Wikipedia evolved, rules changed as to who could edit articles. As experiments in physics have become more sophisticated so too have their rules for collaboration.

In this chapter, I present six arguments: First, collaborations arise for a variety of reasons. Second, collaborations can be usefully categorized as either one-off gatherings which form to take on specific tasks or perpetual efforts that enable us to thrive in a complex world. In both cases, people learn within the collaboration. Third, collaborative architectures matter and influence the success of a collaboration. Participants with requisite knowledge and skills to accomplish a task may be necessary, but this alone is not sufficient to ensure success. Poorly designed architectures can hinder and even undermine a collaboration. Fourth, the space of potential collaborative architectures is enormous, so large in fact to preclude a full taxonomy. As a result, collaborative architectures cannot simply be chosen from a list. This huge space of possibilities means that bespoke collaborative architectures will be typical and that when a collaborative architecture needs to evolve, participants will have many choices.

Fifth, collaborative architectures vary in how they constrain and define acceptable behaviors and norms; this implies that culture, both organizational and geographic, often matters more than might be thought. Finally, within any collaboration, two types of adaptation typically take place: people learn within a collaboration and the architecture itself evolves in response. The resultant architectural adaptation can be formal, with new roles and rules being created, or informal, with initial norms being relaxed or new norms emerging. Therefore, we should view collaborations as coevolutionary processes between behaviors, norms, and the collaborative architecture (Bednar and Page 2018).

From "The Nature and Dynamics of Collaboration,"
edited by Paul F. M. J. Verschure et al. Strüngmann Forum Reports, vol. 33,
Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262548144

The Uses of Collaboration

Groups of all sizes collaborate. People collaborate because groups have a greater capacity than individuals. Collaborative groups are more creative than individuals. Collaborations enlarge and enhance the space of possible solutions to many types of problems (Mulgan 2021). Groups refine causal maps and lead to better solutions by tapping into the diverse knowledge, skills, and perspectives of those who participate in the collaboration (Hong and Page 2001). Collaborations create new functionalities (sometimes planned, sometimes emergent) and encourage specialization, which in turn increases the rate of learning. In the ideal case, as Kant (2008) suggested, the many minds in a collaboration form a singular cognition, resulting in a collective intelligence.

The functions and uses of collaborations are far reaching (Enyedy and Stevens 2014). Collaborations can be used to build or create, to produce knowledge or determine a set of rules, to solve a problem or meet an obligation, to respond to or prepare for an event (Majchrzak et al. 2007). Collaborations can be used to manage common pool resources or to allocate efforts to produce public goods in pursuit of specific goals, such as better managing a school system, reducing traffic, building a more robust power grid, or developing a cure for cancer. Collaborations can also form for the purpose of building a community or to create joy.

Successful collaborations often do both: they achieve a better end and establish community. A book club adds knowledge and understanding among its members and builds friendships and trust between them. Large-scale collaborations, such as those involved in reducing the effects of climate change, strive toward specific goals while building international alliances and personal friendships.

One-Shot and Perpetual Collaborations

Insights into the contributions of collaborations and how they work can be differentiated by viewing them either as one-shot or perpetual collaborations. One-shot collaborations address a fixed problem or undertake a specific task. They do things like solve mathematical theorems, design bridges and hotels, and or write organizational mission statements. Ad hoc committees, such as those that choose company presidents or academic deans, are one-shot collaborations. Not all such collaborations necessarily produce an immediate final output or product. They may lay the groundwork for future decisions or increase trust or understanding among members (Wood and Gray 1991).

Perpetual collaborations differ. These collaborations have lives of their own and can be thought of as adaptive entities. These collaborations, through their participants, iteratively sense, experiment, learn, and adapt to thrive or, in some cases, just to survive. In today's hyperconnected, information-abundant,

From "The Nature and Dynamics of Collaboration,"
edited by Paul F. M. J. Verschure et al. Strüngmann Forum Reports, vol. 33,
Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262548144

and fast-moving world, most ongoing enterprises (e.g., law firm, nonprofit environmental group, community fire department) can be interpreted as engaging in perpetual collaboration within its organizational boundaries.

The mindsets involved in one-shot and perpetual collaborations differ in two important respects. In one-shot collaborations, participants aim to gather relevant information and do deep analytical and empirical dives. For this, participants need to trust one another and communicate effectively, but long-term community building is not necessarily a specific goal. A legal settlement, a classic one-shot collaboration, need not at all be friendly.

In perpetual collaborations, by contrast, deep dives are often less productive because the environment is changing. The collaboration must instead work on developing the capacity to adapt quickly. Perpetual collaborations are like kayakers maneuvering in white water (Pendleton-Jullian and Seely Brown 2018): information flows by at an increasing rate. The problems that the kayakers confront today (e.g., the flow and particular turns of the river) will only weakly resemble those which may be confronted tomorrow, let alone next week. Thus, the architecture that underpins a perpetual collaboration must be less formal and more adaptive. Given this need to adapt, increasing and then maintaining trust and building communication skills become central goals. In sum, within a perpetual collaboration, building the collaborative community takes priority, whereas one-shot collaborations focus more on the task at hand.

Thinking Two-by-Two: Why Collaborative Architectures Matter

Collaborations occur within architectures, and those architectures must be suited to the task. Thus, a collaboration between jazz musicians in real time during an impromptu concert will be much less formal than a collaboration among surgeons and nurses involved in high-risk brain surgery.

Some linkages between task and architecture are straightforward: On tasks that can be divided or modularized, people work asynchronously and make decisions on their own. On a disjunctive task, where the collective solution equals the value of the best solution, a collaboration can be open and loose. Improv theater troupes encourage interruptions with the hope that someone will produce a huge laugh.

In contrast, on a conjunctive task, where the value of collaboration equals that of the worst performer, or on a highly interdependent task, greater control over decisions and actions is necessary for success. Hence, orchestras have conductors, football teams have quarterbacks, and department stores have managers.

Organizational scholars have created a number of taxonomies that relate structure to performance. Pisano and Verganti (2008) construct a two-by-two collaborative architecture taxonomy based on whether participation is open or closed and whether decisions are made within a hierarchy or decided on by

From "The Nature and Dynamics of Collaboration,"
edited by Paul F. M. J. Verschure et al. Strüngmann Forum Reports, vol. 33,
Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262548144

the participants in a flat decision-making structure (see Figure 14.1). Though intuitive, these distinctions clarify our thinking about types of collaborations and why they look like they do. For instance, the earliest version of Wikipedia was open: anyone could edit an article. The collaboration *among* writers for a president's State of the Union speech is closed: tourists visiting Washington, DC, from Stevens Point, Wisconsin, or Bend, Oregon, simply cannot stop by and offer their ideas. Decision making in most militaries is hierarchical: the highest-ranking officer decides. In most research collaborations, coauthors may all approximately have equal say and reach decisions by consensus. In each of these cases, the architecture "fits" the task.

These two distinctions create four pure types of collaborations as shown in Figure 14.1. A circle/team consists of a selected set of people working together as part of an organization (perhaps scientists or engineers) on a problem chosen by the organization. The organization, or the person or people at its head, makes the important decisions. A consortium differs from a circle/team in that the participants collectively decide on the problems and solutions. The members of a consortium may even have different goals.[1] The bottom row describes collaborations with open participation. Crowdsource captures collaborative platforms, like InnoCentive, in which a central organization posts a problem with prize money attached. Anyone can participate in InnoCentive as a problem solver and earn the prize money, but the organization defines the prize and determines who wins. By contrast, in a community, such as the Linux source code development community, both the problems and solutions are defined by the community.

Each design choice has advantages and disadvantages. Closed collaborations typically rely on paid experts, so quality is less of a concern; however, people with relevant knowledge need to be identified. Open collaborations, by contrast, have access to many ideas and different types of knowledge. For instance, InnoCentive takes problems unsolved by "the experts" and challenges a broader community to provide insights or resolve the problem.

		Decision Making	
		Hierarchical	*Flat*
Participation	*Closed*	• Circle/team • Elite group • Chosen by company	• Consortium • Small group • Closed
	Open	• Crowdsource • Company poses problem • Company decides winner	• Community • Network of people • All can pose problems

Figure 14.1 A two-by-two collaborative architecture taxonomy to manage performance in a collaboration (after Pisano and Verganti 2008).

[1] Under this taxonomy, the Charlie Parker jazz quintet, which included Miles Davis, would be a consortium.

From "The Nature and Dynamics of Collaboration,"
edited by Paul F. M. J. Verschure et al. Strüngmann Forum Reports, vol. 33,
Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262548144

Collaborations using hierarchical or centralized decision rules need to be able to identify when something represents an improvement or a path to an improvement. Flat, consensus-based decisions leverage diverse perspectives at a cost of pulling in opposite directions. Consistent with this framework, flatter scientific research teams are more innovative, whereas hierarchical teams are better at carrying out mainstream research (Xu et al. 2022).

No one of these categories should be considered better than the others. Instead, each should be thought of as best suited to particular tasks. A closed hierarchical architecture (circle/team) might work best for a group of scientists working for a chemical company tasked with finding a solution that satisfies certain properties. A consortium of a group of firms (i.e., closed flat architecture) will perform better than a circle/team within a single firm, when complex domains with increasing knowledge bases and multiple sources of expertise are involved (Powell et al. 1996). When sources of expertise are widely dispersed, the locus of innovation will be found in networks of learning rather than in individual firms (Powell 1990). Crowdsourcing (open hierarchy) works well when knowledge is held broadly and the problem is well specified, as is the case with Kaggle programming challenges. Finally, a community (open flat architecture) works best when neither the problem nor good solutions are well defined. Early-stage startups often rely on flat structures with collective decision making. That makes sense given that they are often operating under high ambiguity.

The Space of Possible Collaborative Architectures

The two-by-two taxonomy illustrates why architecture matters, but it understates the enormity of the space of possible architectures. When designing a collaboration, there are far more decisions to be taken than simply who belongs and how decisions will be made. Collaborative architectures assign roles, create incentives, define interaction structures (synchronous or asynchronous, virtual or in person), design spaces for sharing information, and establish communication protocols and boundaries.

The task of designing a collaborative architecture could be framed as a mechanism design problem (Myerson 2008). Mechanism design assumes an environment that characterizes the information of participants, their preferences, and relevant details about the context. In the case of an economy, it would include initial allocations of resources and current technologies. A mechanism consists of a message space, which could be prices, votes, or even words, along with a mapping from sets of messages to outcomes.

Though a powerful framework, mechanism design relies on strong assumptions and abstracts away from the types of details considered here. Nevertheless, it introduces important insights, chief among them that any mechanism (a collaboration or otherwise) may or may not create incentives for truth telling or

From "The Nature and Dynamics of Collaboration,"
edited by Paul F. M. J. Verschure et al. Strüngmann Forum Reports, vol. 33,
Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262548144

cooperating. From the perspective of mechanism design, an ongoing collaboration can be thought of as a repeated game in which punishment strategies can encourage good behavior.

In mechanism design, the objective (the social welfare function) is known. In a collaboration, the objective may itself be an output of the collaboration: What is to be done? As the objective crystallizes, the architecture may change. Thus, collaborations also need rules or processes for changing the architecture. This may involve adding or removing members, changing decision-making rules, or altering communication protocols (Arrow 1974). We can, therefore, distinguish architectures as self-organized, organizationally enabled, or organizationally structured (Engeström et al. 2015). Some collaborations begin with a loose set of rules and let structures and protocols emerge. Other architectures hardwire the connections, communication rules, and behavioral protocols (Fjeldstad et al. 2012).

In an architecture, each feature requires thoughtful construction on its own. Getting the incentives right is not simply a matter of increasing compensation. In complex collaborations with multiple objectives, monetary incentives will not be enough (Weiss and Hughes 2005). Similarly, we might think that consensus works best as a decision-making rule and often that is true. It can even hold when participants vary in their expertise, such as with a group of doctors and patients working to improve safety practices that relied on consensus (Trier et al. 2015). However, in the case of a scientific committee evaluating the societal dangers of more than a dozen drugs on more than a dozen criteria, reaching a consensus would take a prohibitive amount of time. This sort of collaboration requires a formal voting or rating procedure (Nutt et al. 2010).

Writers, Research, Open-Source Programming

The features of a collaboration are also interdependent, as can be seen through three examples of collaboration: a sitcom writers' room, a scientific research team, and an open-source programming collaboration.

The roles in a writer's room include lead writer and staff writer. The interaction structure might include everyone in the room at the same time. In this case, the message space would likely consist of the words in a common script. The lead writer might present that initial script, and other writers would then suggest changes. The decision-making mechanism could be dictatorial, with the lead writer deciding. However, given that the lead writer wants people to be comfortable sharing ideas, such a rule might not work well. More likely, the lead writer would be more inclusive and permit changes preferred by a majority in the room. Alternatively, the lead writer might require a supermajority, say five of seven writers. Incentives, in this case, would be both direct (i.e., the writers' pay might be tied to the success of the show) and indirect (i.e., staff writers who consistently make good suggestions would improve their reputation within the room, leading to other opportunities). Note the interdependence

From "The Nature and Dynamics of Collaboration,"
edited by Paul F. M. J. Verschure et al. Strüngmann Forum Reports, vol. 33,
Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262548144

of the architecture. The common script requires a decision-making process at each step. That would not be true for other collaborations.

In a scientific collaboration set up to study the effects of stress on blood pressure, participants might be a select group of scientists whose roles are tailored to their expertise. Participation would be closed (i.e., by invitation only). Some members who design the experiments, for instance, would interact as a fully connected group. This subgroup might meet in person, use a white board to design experiments, and exchange experimental designs per email or via a Slack channel that includes critiques and commentary. Decisions would occur only sporadically. They could be made by majority rule or be hierarchically based (e.g., by the most senior scientist or, in funded research, by the grant's principal investigator).

Unlike in the writer's room, in the scientific collaboration, participants can explore alternatives in depth. There is no common script. Scientific collaborations involve groups of experts with specific skills and tools. Some members of a collaboration might focus on the analysis of data. These scientists might interact with only one member of the experimental design group. Initially, they might work in isolation to produce independent analyses before coming together as a group to compare their findings, at which point, they might correct and refine their respective analyses. Ultimate choices may be made by consensus. Their final product might be a single analysis along with a collection of redundancy checks. Their analyses (another form of message) might then be sent to the lead author (another role), responsible for drafting the academic paper or scientific report. Upon completing that report, the lead author might send the report to all the scientists for commentary. These scientists might be able to make recommendations, which the lead author can adopt or ignore, or raise objections, which the lead author is obliged to correct.

If the lead author disagrees with the objection, a meeting or vote of all the scientists may be required to adjudicate. The incentives for the scientists should align, though not perfectly. While all of the scientists want the research project to be well received (e.g., to earn citations and future funding), they may want different parts of the contribution to be highlighted in the final report. Some may care more about the scientific contribution, while others may be more concerned about the methodological advances. In addition, junior scientists may be as concerned with learning new techniques or making professional connections as with the research itself.

As a final example, let us consider a collaboration for open-source programming to build a software program (e.g., Linux algorithm). This type of open collaboration allows anyone to participate. The only formal roles would be the project managers. Messages might take the form of computer code as well as queries and comments within self-organized Slack channels. Decisions would likely be made by committee or consensus, depending on their centrality to the project. The incentives for such a collaboration may be nonmonetary. People may put in effort to learn, to build reputations, to compare their skills

From "The Nature and Dynamics of Collaboration,"
edited by Paul F. M. J. Verschure et al. Strüngmann Forum Reports, vol. 33,
Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262548144

against others, to contribute to scientific advancement, or some combination of these reasons.

These three examples demonstrate the breadth of possible collaborative architectures, the interdependence of their features, and make obvious why a complete taxonomy is out of reach, as it would have to include everything from the garbage can model (Cohen et al. 1972) to pure hierarchies. To be of use, such a taxonomy would also have to explain the contexts in which each architecture performs well. An architecture may seem well suited to a large set of domains but prove to be effective only under a narrow set of circumstances (Bendor et al. 2001).

These three examples also reveal a link between collaborative architectures and behaviors and norms. Writers and open-source programmers take more of a doing, using, interacting approach and would likely develop a looser set of norms, whereas scientific collaborations might produce more regimented and formal behaviors (Jensen et al. 2007).

Collaborative Architectures in Context: Culture Matters

Collaborative architectures must not only be structured relative to task or purpose, they must also be attuned to the cultures within which they are embedded. In a high trust, homogeneous culture, rules for equal speaking time or equal access to resources used in the collaboration may be unnecessary. In a low trust, diverse culture, rules for collaboration may need to be explicit to ensure fairness. A diverse culture may also require more elaborate communication protocols to ensure information and knowledge transfer.

Cultures differ according to their *tightness* and *looseness* (Gelfand et al. 2011). Loose cultures have weak social norms; people break rules (cut in line, interrupt one another, disrespect authority). Tight cultures have strong social norms; deviant behavior is not tolerated.

The same collaborative architecture can perform differently in tight and loose cultures. Although some might think that tight cultures perform better on all tasks, this need not be the case. People in a tight culture might find themselves overwhelmed by emails and meetings (Polzer et al. 2018). Such "over collaboration" can produce fatigue and result in poor performance (Polzer and DeFilippis 2020). Therefore, in a tight culture, the collaborative architecture may have to restrict communication and promote free thinking. By contrast, in a loose culture, the architecture may oblige reporting and monitoring of the reporting. What works in Singapore (a tight culture) might not function effectively in the Netherlands (a loose one). What works on Wall Street (a tight culture) might fail miserably on Sand Hill Road (a loose one).

The *competing values framework* (Figure 14.2) links culture to organizational type (Cameron and Quinn 1992) and distinguishes organizations according to their focus (inward or outward) and structure (stable or flexible). In this

From "The Nature and Dynamics of Collaboration,"
edited by Paul F. M. J. Verschure et al. Strüngmann Forum Reports, vol. 33,
Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262548144

		Focus	
		Inward	*Outward*
Structure	*Flexible*	Hierarchy: • Structured • Process and efficiency focused	Market: • Competitive • Achievement focused
	Stable	Clan: • Cohesive • Nurturing • Family structure	Adhocracy: • Adaptive • Risk-taking • Innovative

Figure 14.2 Competing values framework, linking culture to organizational type (after Cameron and Quinn 1992).

framework, structure and culture interact. An inwardly focused organization with a stable structure will result in a process- and efficiency-minded culture. At the other extreme, an outwardly focused, flexible structure produces a culture that is adaptive and not adverse to risk taking.

No single culture possesses the ideal prescription to address all unknown emergent challenges. Collaborative cultures must be adaptive, but they also benefit from consistency. Thus, efforts to encourage a hierarchical and scientific culture, while also emphasizing failing fast and learning, can lead to noncomplementary innovations (Haus-Reve et al. 2019).

Evolving Collaborative Human Capital

How effectively a collaborative architecture performs depends in part on culture; in particular, on norms, beliefs, and behaviors. Do people, for example, follow rules? Do they have high levels of trust? In perpetual collaborations, culture and collaborative architectures coevolve. Through collaboration, people build human capital which in turn makes them better collaborators.

Lansing and Cox (2019) provide a compelling example of the coevolution of architecture and collaborative human capital. They describe how Balinese rice farmers choose harvest schedules that balance the benefits of simultaneous harvests, which decreases damage from pests, against the increased cost of trying to engage every farmer at the same time. This collaboration among farmers occurs in nonhierarchical organizations called *subaks.* These *subaks* allocate water flows through a series of water temples, a collaborative architecture that has emerged over centuries. Not only has the architecture evolved, so too have the behaviors and norms that support the system.

Evidence for the development of *collaborative human capital*, the development of tools that enable better collaborative performance among a given set of participants, exists in a variety of settings—from basketball (Maymin et al. 2013) to research and development alliances (Sampson 2005). Whether the evolution

From "The Nature and Dynamics of Collaboration,"
edited by Paul F. M. J. Verschure et al. Strüngmann Forum Reports, vol. 33,
Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262548144

of collaborative human capital can transfer from one domain to another likely depends on the similarities of the collaborative architectures and cultures.

Summary

Collaborative architectures structure how people collaborate: Who belongs? What roles will they assume? How are decisions made? How does communication occur? How is good behavior rewarded and bad behavior punished? How can the rules be changed? The type of architecture one chooses is dependent on the goals and tasks to be accomplished. Innovative collaborations benefit from flatter, nonhierarchical structures whereas more technical, scientific collaborations often require precise communication protocols.

Choice of architecture must also be mindful of culture. A hierarchically organized collaboration with strict rules for who speaks, when, and on what topics may produce good outcomes in a tight culture but perform poorly in a loose one. Alternatively, an open, flat collaborative structure may not work well in a tight culture, even when the goal is to produce innovative ideas.

Collaborative architectures can take an enormous number of forms, and this plethora of possibilities offers potential as well as risks. Any given collaborative architecture might perform spectacularly in certain contexts yet fail miserably in others. Thus, astute designers of collaborations rely on frameworks and logic when contemplating structural elements. They also build in rules for adapting the architecture over time. These adaptations can be made to improve performance toward an existing goal, to steer the collaboration toward new goals, or to account for the changes in skills and knowledge of the collaborators. Guiding the evolution of the collaborative architecture is itself a collaborative process. As participants become better collaborators within an architecture, they may also become better at adapting that architecture.

From "The Nature and Dynamics of Collaboration,"
edited by Paul F. M. J. Verschure et al. Strüngmann Forum Reports, vol. 33,
Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262548144

From “The Nature and Dynamics of Collaboration,”
edited by Paul F. M. J. Verschure et al. Strüngmann Forum Reports, vol. 33,
Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262548144