STABLE OR ROBUST? WHAT'S THE DIFFERENCE?

Erica Jen
Santa Fe Institute
Santa Fe NM
erica@santafe.edu

December 17, 2002

1 Abstract

Exploring the difference between "stable" and "robust" touches on essentially every aspect of what we instinctively find interesting about robustness in natural, engineering, and social systems. It is argued here that robustness is a measure of feature persistence in systems that compels us to focus on perturbations, and often *assemblages* of perturbations, qualitatively different in nature from those addressed by stability theory. Moreover, to address feature persistence under these sorts of perturbations, we are naturally led to study issues including: the coupling of dynamics with organizational architecture, implicit assumptions of the environment, the role of a system's evolutionary history in determining its current state and thereby its future state, the sense in which robustness characterizes the fitness of the set of "strategic options" open to the system; the capability of the system to switch among multiple functionalities; and the incorporation of mechanisms for learning, problem-solving, and creativity.

2 Introduction

"What's the difference between stable and robust?" It's the first question that comes to mind, especially for researchers who work with quantitative models or mathematical theories. Answering the question isn't made any easier by the fact that "robustness" has multiple, sometimes conflicting, interpretations—only a few of which can be stated with any rigor. (For a list of working definitions of robustness, see the Santa Fe Institute robustness web site at http://discuss.santafe.edu/robustness.)

But in fact the question of the difference between "stable" and "robust" touches on essentially every aspect of what we instinctively find interesting about robustness. It's worth trying to answer the question even if the answers are wrong (or, as is more likely, even if the answers are too vague to be either right or wrong).

It may help to ease into the topic by asking, in order,

"What is stability?"
"What do stability and robustness have in common?"
"What is robustness beyond stability?"

The concept of stability is an old one that derives from celestial mechanics and in particular the study of the stability of the solar system. A readable treatment is provided by Wiggins [35]. Definitions will be paraphrased here for the sake of establishing some basic language.

Loosely speaking, a solution (meaning an equilibrium state) of a dynamical system is said to be *stable* if small perturbations to the solution result in a new solution that stays "close" to the original solution for all time. Perturbations can be viewed as small differences effected in the actual state of the system: the crux of stability is that these differences remain small for all time.

A dynamical system is said to be *structurally stable* if small perturbations to the system itself result in a new dynamical system with qualitatively the same dynamics. Perturbations of this sort might take the form of changes in the external parameters of the system itself, for example. Structural stability requires that certain dynamical features of the system, such as orbit structure, are preserved, and that no qualitatively new features emerge. Diacu and Holmes [9] give the example of flow on the surface of a river to illustrate the notion of structural stability. Assuming that the flow depends on an external parameter, such as wind speed, and ignoring other factors, the flow is structurally stable if small changes in wind speed do not qualitatively change the dynamics of the flow; for example, do not produce a new structure such as an eddy.

As for the commonalities between stability and robustness,

1. most if not all communities would agree both concepts are defined for *specified features* of a given system, with *specified perturbations* being applied to the system. It makes no sense to speak of a system being either stable or robust without first specifying both the feature and the perturbations of interest.
2. both stability and robustness are concerned with the *persistence*, or lack thereof, of the specified features under the specified perturbations. Persistence therefore can be seen as evidence of either stability or robustness.

So what's the difference between stability and robustness? It'll be argued here that robustness is broader than stability in two respects. Firstly, robustness addresses behavior in a more varied class of

- systems;
- perturbations applied to the system of interest;
- features whose persistence under perturbations is to be studied.

Second, robustness leads naturally to questions that lie outside the purview of stability theory, including

- organizational architecture of the system of interest;
- interplay between organization and dynamics;
- relation to evolvability in the past and future;
- costs and benefits of robustness;
- ability of the system to switch among multiple functionalities;
- anticipation of multiple perturbations in multiple dimensions;
- notions of function, creativity, intentionality, and identity.

In what follows, arguments will be given to support the above two points. All arguments are based primarily on plausibility, and are in urgent need of both empirical and theoretical elaboration.

3 Contexts for Stability Versus Robustness

It's easy to list examples of systems, features, and perturbations for which the language and framework of traditional stability theory—whether addressing the stability of states or the structural stability of the system—seem inadequate. The table below includes examples for which stability theory is entirely appropriate, and others that arguably call for some notion of robustness different from stability. It'll be left as an exercise to the reader to say which is which.

System	Feature of Interest	Perturbation
earth's atmosphere	temperature	increase in fluorides
rangelands	biomass	change in grazing policies
laptop	software performance	incorrectly entered data
laptop	software performance	disk crash
bacterial chemotaxis	adaptation precision	change in protein concentrations
bacterial chemotaxis	adaptation precision	bacterial mutation
human immune system	antibody response	new virus
human immune system	antibody response	autoimmune disorder
U.S. political system	perceived legitimacy	demographic changes
U.S. political system	perceived legitimacy	economic depression
religions	popularity	modernity
footbinding[1]	longevity	change in status of women
automotive market [2]	identity of Volkswagen "Bug"	changes in design

What characterizes the contexts in which robustness captures some aspect of a system different from those described by stability theory?

---

[2] Suggested by Sam Bowles.

[2] Suggested by Josh Epstein.

The first observation is that robustness is a measure of feature persistence for systems, or for features of systems, that are difficult to quantify, or to parametrize (i.e., to describe the dependence on quantitative variables); and with which it is therefore difficult to associate a metric or norm. The differences among various auction designs and their robustness to collusion, say, is difficult to describe from a traditional stability or structural stability perspective. It is easy even in the context of traditional dynamical systems to define qualitative features (equivalence classes of attractors, details of phase transitions, for example) that would be a stretch for stability analysis [12].

Second, robustness is a measure of feature persistence in systems where the perturbations to be considered are not fluctuations in external inputs or internal system parameters, but instead represent changes in system composition, system topology, or in the fundamental assumptions regarding the environment in which the system operates. Morever, as pointed out by David Krakauer [18], it is typical in stability theory to postulate a single perturbation; from the robustness perspective it is often ineluctably necessary to consider instead multiple perturbations in multiple dimensions. A biological signalling pathway for example may be robust to an entire assemblage of perturbations including not only fluctuations in molecular concentrations but also the "knocking-out" of an array of different genes all of which *prima facie* appear essential in different ways to the functioning of the pathway.

Robustness moreover is especially appropriate for systems whose behavior results from the interplay of dynamics with a definite organizational architecture. Examples of organizational architectures include those based on modularity, redundancy, degeneracy, or hierarchy, among other possibilities, together with the linkages among organizational units. The redundancy and degeneracy of the genetic code [17], the functional modularity of ecosystems [19], the hierarchical nature of regulatory regimes [22]—these are examples of organizational features not easily represented in a stability framework. Even more importantly, these organizational features are in many systems spliced together into what social scientists term "heterarchies" [28]; namely, interconnected, overlapping, often hierarchical networks with individual components simultaneously belonging to and acting in multiple networks, and with the overall dynamics of the system both emerging and governing the interactions of these networks. Human societies in which individuals act simultaneously as members of numerous networks— familial, political, economic, professional, among others—are one example of heterarchies, and signalling pathways in biological organisms are another, but in fact the paradigm is a powerful one with relevance to many natural, engineering, and social contexts.

Note that robustness is meaningful for heterarchical and hierarchical systems only when accompanied by specification of the "level" of the system being so characterized. In an ecosystem, for example, the individual constituent species may be robust with regard to certain disturbances, and interact in such a way as to give rise to a similarly robust aggregate. Even with species that are themselves robust, however, the ecosystem as a whole may not be robust. Even better, species that are not themselves robust can undoubtedly interact so as to create a robust aggregate. In other words, presence or absence of robustness at one level does not imply presence or absence at another level, and perhaps the most interesting cases are those in which the interconnections among components not themselves robust give rise to robustness at the aggregate level [20, 34].

Implicit in the above is the idea that robustness typically applies to what, for lack of better terminology, are often called "complex adaptive systems." As John Holland points out, "Usually we don't care about the robustness of a rock." In many of these cases, robustness may be interpreted as an index of the relative strengths and weaknesses— what might also be called the "fitness"—of the set of "strategic options" that either have been designed top-down or have emerged bottom-up for the system. The options available to the system serve in other words as a "strategy" for how to respond to perturbations.

The concept of robustness as applying to systems with strategic options is useful in unifying two ostensibly different interpretations of the term. Robustness is often thought of as reflecting the ability of a system to withstand perturbations in *structure* without change in *function*—in biological contexts, this is sometimes called "mutational robustness," and as argued above may be seen as measuring the fitness of a strategy that has either emerged, or has been selected, for responding to insult or uncertainty. However the dual interpretation of robustness is equally valid; namely, robustness may be seen as measuring the effectiveness of a system's ability to switch among multiple strategic options. Robustness in this sense reflects the system's ability to perform multiple functionalities as needed *without* change in structure—this might be called "phenotypical plasticity." As an example of the second type of robustness, simple invertebrate neuronal networks can be seen as heterarchies with the ability to reconfigure themselves and to modify the intrinsic membrane properties of their constituent neurons [23]—the result of this neural "strategy" is a form of robustness that enables a single network to perform, and to switch among, multiple functional tasks.

Strategy is associated with systems that are acting to some purpose, and indeed, robustness is usually ascribed to systems perceived has having a well-defined function or performance measure. Often the function or performance measure is seen as **the** central feature of the system. The concept of function is however problematical in almost all cases. Identifying function with a physical process such as a fluid flow is certainly in the mind of the beholder, and it's not clear that it brings any insight whatsoever. In the case of ecosystems, the concept of function is usually replaced by that of nutrient cycling, productivity, or other "functionals" of the system that may or may not correspond to intuitive ideas of ecological resilience.

As for social systems such as stock markets or religious institutions, the assignment of function may arguably be more natural but equally subjective, in part because such systems acquire in the course of their development multiple functions both consistent and inconsistent. Moreover, as Levin [21] has pointed out for ecosystems, "robustness" for social institutions often becomes synonymous with rigidity. The *survival* of social institutions such as firms or bureaucracies or governments may sometimes emerge as their primary "function"; adaptation and evolution in these cases largely represent attempts to maintain legitimacy rather than improvements in the institutional function as it was originally defined [10]. The relevance of robustness to the long-term future of social systems will be discussed in [13].

Even for engineering or computational systems, however, the notion of "function" has pitfalls. Intentions of the designers notwithstanding, the systems may possess multiple functionalities in part as a function of its heterarchical nature, and as in the example of the neuronal networks, it may be the ability to switch among tasks that is the true feature of interest. Certainly the possibility cannot be excluded that the systems may develop new functionalities unanticipated in their design. "Function" may therefore not be as important a distinction between robustness and stability as at first it appears, but remains nevertheless an issue that requires careful attention.

4 Robustness Beyond Stability

Many questions that sit uneasily within a stability framework arise naturally in a study of robustness.

It's argued above that robustness is a concept appropriate to measuring feature persistence in certain contexts; namely, systems where the features of interest are difficult to parametrize, where the perturbations represent significant changes either in system architecture or in the assumptions built into the system through history or design, or where the system behavior is generated through adaptive dynamics coupled to strong organizational architecture. The study of robustness then naturally prompts questions relating to organization, the role of history, the implications for the future, and the anticipation of insults, along with other questions even more difficult to formulate relating to creativity, intentionality, and identity.

4.1 Interplay between Dynamics and Organization

What is the interplay between the dynamics and organizational architecture of robust systems, and how does the architecture both facilitate and constrain the dynamics by which robustness may be achieved? Hartman et al [11] argue, for example, that it is this interplay that permits "living systems to maintain phenotypic stability in the face of a great variety of perturbations arising from environmental changes, stochastic events, and genetic variation." One simple example of a coupling between dynamics and organization in a computational context is the use of a a majority voting rule to resolve conflicting inputs and to provide an error-correcting capability [31, 3]. In biological contexts, a growing literature [11, 15, 17] provides empirical evidence and theoretical conjectures on more sophisticated dynamics—including for example neutral evolution, and positive or negative feedback—that serves a similar purpose. Much of this work describes the diverse and intricate array of molecular and cellular mechanisms that permit the accumulation of genetic variation while buffering the organism against deleterious phenotypic consequences, or that faciliate higher-level mechanisms to identify and to correct such variations as needed.

Arnold and Voigt [30] provide another perspective on the interplay between dynamics and architecture. They argue that in the context of directed evolution, organizational structure can be exploited to accelerate the discovery of proteins with novel or prespecified catalytic properties. In particular, the approach is based on separating the units that modulate function from those that maintain structure. As they find, "by making specificity-determining units structurally tolerant, the space of possible functions can be explored quickly."

The role of organizational architecture in generating robustness is obscure in part because the origins of organization are themselves murky. It may be that the converse question is in fact the correct one for many systems, especially in engineering and computational context; namely, "What is the role of robustness in generating organizational architecture?" Carlson and Doyle, for instance, argue that much of the complexity in sophisticated engineering systems stems not from the specifications for functionality, but from the exigencies of robustness [7]. They argue that in traditional engineering design, regulatory mechanisms for robustness are typically superimposed after the fact on the mechanisms for functionality, and that Rube Goldberg prevails as a consequence.

A different view of the interplay between organizational architecture and robustness emerges from the study of certain hierarchical systems. As pointed out in the previous section, the discussion of robustness for such systems has meaning only when the level of the system is clearly identified. Robustness may exist on the level of the individual components, or on an intermediate level, or on the level of the whole, or not at all. Robustness on one level need not imply robustness on any other level. Conversely, robustness at one level may—through processes seredipitous or otherwise—confer robustness at another level (see work by Ancel and Fontana [1] on mutational and environmental robustness in RNA evolution, for example, and by Marder [24] indicating the role of cellular mechanisms for plasticity in ensuring higher-level neuronal and circuit stability). Are there what Ancel and Fontana call "congruence principles" for translating robustness on one level to robustness at another level?

4.2 History and Future

The role of a system's history, and the implications for its future, represent a second set of questions that are stimulated by the study of robustness. In evolutionary and developmental terms, the specific nature of a system's robustness may both reflect the legacy of its history, and constrain the realizations possible in its future [27].

Intuition tells us for example that there are tradeoffs between robustness and evolvability. Robustness loosely speaking may be seen as insensitivity to external and internal perturbations. Evolvability on the other hand requires that entities alter their structure or function so as to adapt to changing circumstances.

Kirschner and Gerhart[15] argue however that the existence of redundancy and drift can introduce "useful" variability into a system. If the system can be protected from lethal mutations, the accumulation of variability may permit the system to move to a state within the same neutral network—sets of systems that are "genotypically" different although "phenotypically" equivalent—such that fewer subsequent mutations are needed to effect a major innovation. This form of robustness thus exploits the combination of redundancy and the dynamics of drift and neutral evolution in order to increase evolvability. The precise distribution of the neutral network within the space of all systems will determine not only the accessibility of the system from historical antecedents, but also the course and the robustness associated with future innovations [29].

What is the relation between robustness and historical persistence (mimicking those who ask the relation between conservation throughout evolutionary history and the adaptedness of genes)? The question is relevant to phenomena throughout the natural, engineering, and social spheres, but it's the rare case where even tentative conclusions can be drawn. One such case is provided by mammalian sensory processing systems, for which it can be inferred from empirical data that the systems that are oldest in evolutionary terms are in some sense the most robust [36]. The sensory system for taste, for example, can regenerate itself after destruction of the taste buds, and the olfactory system can similarly recover from loss of the chemoreceptors. However, the taste system—which predates the olfactory system in evolutionary terms—apparently can regenerate itself after loss of enervating nerves, whereas the olfactory system cannot.

Krakauer [18] argues in this context that persistence may result from any number of reasons—including constancy of environment, or constraints developmental and evolutionary—that do not necessarily imply robust design. Moreover feature persistence in a population may merely reflect the fact that the feature is deleterious in individuals and hence individuals with the feature do not survive. Robustness in that case can only be seen as an after-the-fact property of the population that is generated as a byproduct of the constraints on individuals. In general, what are the contexts in which persistence may be taken as evidence of robustness, and what are the mechanisms by which those cases of persistence are realized?

As alluded to above, perhaps the central question in the relation between robustness and history is to distinguish between robustness as a byproduct versus the direct object of selection in either natural or directed evolution. In engineering systems, robustness typically is seen as a direct object of selection (whether part of the initial design or an after-the-fact consideration). This view of robustness may however be oversimplified: Wagner et al [33] for example describe a neural net algorithm that develops a high degree of functional modularity and robustness as a byproduct of its learning to solve a complex problem in robotic control.

In biological examples, the question of distinguishing between byproduct and direct object is even more murky. Have biological systems been expressly selected for robustness [32]? Or is robustness merely a characteristic of systems that contribute directly to high organismal fitness and that in addition survive successive rounds of selection due to their tolerance of high variability and their ability to generate phenotypic novelty? Has the chemotaxis pathway been selected for its robustness, or is robustness a byproduct of a highly fit functional design [2]? To what extent is the robustness of the pathway a feature of importance to the organism? What is the difference between robustness and regulation in this case?

4.3 Robustness as Fitness of Strategic Options

If robustness can be viewed as characterizing the fitness of the "strategic options" open to a system (whether the options have been selected for or emerged), then there are likely to be important system-wide consequences associated with that set of options. One such consequence is the balance of costs and benefits of particular form of robustness. For example, in the auditory system there may be energetic costs to be traded off against the performance benefits of redundant processing of signals by arrays of hair cells. As a general principle, Carlson and Doyle [4] conjecture the existence of strict "conservation laws" for robust systems that require that high tolerance to certain insults be accompanied necessarily by low tolerance to other insults. An aircraft can either be highly manueverable or be extremely resistant to anti-aircraft fire, but probably not both.

Other consequences might relate to the number and type of options opened up or closed off by the particular form of robustness. The political regime in Renaissance Florence established by Cosimo de'Medici (1389-1464), for example, is analyzed by Padgett [25] as deriving its robustness from a strategy of flexible opportunism that permits actions to be "interpreted coherently from multiple perspectives simultaneously," with the consequence of maintaining what Padgett calls "discretionary options across unforeseeable futures in the face of hostile attempts by others to narrow those options."

Moreover, as a different type of constraint, systems that are robust often are required to maintain function while exploring new functionality. Procedures to upgrade networking protocols on the Internet, for example, must implemented without interrupting functionality. Software engineers refer to this principle as "online management," and Walter Fontana likens it to the "need to fix a boat out on the water rather than on dry dock." What are the implications for the evolvability of robust systems?

4.4 Identification of and Response to Insults

The use of robustness as a design principle raises a deep set of questions as to the nature of, and the response to, assemblages of insults previously unencountered and in a real sense unforeseeable. What design principles should guide the construction of systems for which there exists an infinite number of possible insults to which the system may be subjected? The possibility of using joint probability distributions to estimate the likelihoods and consequences of failure is fairly dim here. What other tools can be developed to endow a system with "open-ended robustness," and what would such a system look like?

One comment in this regard is that robustness of this nature, if it exists, would share some characteristics with the higher-level cognitive processes of the brain (which is of course in the processes of development and learning a quintessentially robust system). The conjecture is that "open-ended robustness" would depend on the performance of the system in dimensions such as induction and deduction, the emergence of innovation, and creative problem-solving.

As a particular example of the challenge of modeling insults, what is the difference between designing robustness against "purposeless" versus "purposeful" perturbations or attacks? To first order, stability theory can be said to address the consequences of perturbations that lack intentionality. By contrast, as pointed out by Schneider [26] and Kearns [16], the robustness of computer network security systems is an example—as is the rule of Cosimo de'Medici mentioned above—in which it is necessary to posit instead the existence of attackers intimately familiar with the vulnerabilities of the specific system, and in possession of the expertise and resources needed to mount a coordinated attack explicitly designed to cripple or to destroy. Robustness to this form of attack clearly calls for design that includes the ability to learn, to anticipate, and to innovate.

4.5 Robustness and Identity

Finally, in some important contexts the feature of a system that is robust to disturbances is the *identity* of the system itself. (See the Volkswagen Bug example in Table 1, as well as examples from physical systems capable of self-assembly and self-repair [8].) In such cases, does there exist an instruction set or memory that permits the system to preserve its identity even under severe disruption? What are the mechanisms of repair and of self-maintenance? Why does an observer choose to perceive a system as robust despite its perhaps having undergone fundamental changes either structurally or functionally? The issue of identity—not one often highlighted in the natural sciences—underlies the study of robustness, and raises questions for which many disciplines lack even the language in which to pose them.

5 Summary

In its weakest form, the argument for robustness as different from stability can be stated as follows:

> Robustness is an approach to feature persistence in systems for which we do not have the mathematical tools to use the approaches of stability theory. The problem could in some cases be reformulated as one of stability theory, but only in a formal sense that would bring little in the way of new insight or control methodologies.

In stronger form, the argument can be stated as:

> Robustness is an approach to feature persistence in systems that compels us to focus on perturbations, and assemblages of perturbations, to the system different from those considered in the design of the system, or from those encountered in its prior history. To address feature persistence under these sorts of perturbations, we are naturally led to study the coupling of dynamics with organizational architecture; implicit rather than explicit assumptions about the environment; the role of a system's evolutionary history in determining its current state and thereby its future state; the sense in which robustness characterizes the fitness of the set of "strategic options" open to the system; the intentionality P of insults directed at, and the responses generated by, the system; and the incorporation of mechanisms for learning, innovation, and creative problem-solving.

The above "Strong Form" of the thesis might at first glance appear to rule out applicability to essentially any real system in any but a metaphorical sense. "Strategic options" for biological systems? And yet the interpretation of biological systems acting "on their own behalf" has proved useful in several contexts [6, 14]. As for physical systems—which ostensibly lack organizational architecture along the lines typical of biological systems—it is not impossible that the insights from studying the robustness of hierarchical systems, say, will assist in the understanding of physical processes across multiple scales. (Note that the converse is also true, see for example [4]). Certainly with respect to engineering and computational systems, the evolutionary dynamics (both past and future) of these systems represents a topic with enormous potential for illuminating principles of robust design.

But the proof is in the pudding, or as is said in Chinese, what is needed is to "tou bi cong rong" ("throw down the pen to join the army"). Important distinctions between "robust" and "stable" notwithstanding, the study of robustness as a design principle of natural, engineering, and social systems will become meaningful only if its use in some specific context results in an interesting insight that couldn't have been gotten otherwise. Stay tuned.

Acknowledgments. This essay is the outgrowth of conversations with many participants in the SFI program on robustness. Discussions with Homayoun Bagheri, Sam Bowles, Leo Buss, Josh Epstein, Walter Fontana, Juris Hartmanis, Leo Kadanoff, Michael Kearns, Simon Levin, John Padgett, John Pepper, Eric Smith, Erik van Nimwegen, Gunter Wagner, and Brian Walker have been especially helpful. David Krakauer, John Holland, Steve Carpenter, and George Zweig have influenced the thinking of this essay. I thank Eve Marder for the opportunity to work in her neurophysiology laboratory at Brandeis on the robustness of crustacean stomatogastric systems. The work has been supported by core grants from SFI together with funds from the David and Lucile Packard Foundation.

References

[1] Ancel, L.W. and Fontana, W. "Plasticity, evolvability, and modularity in RNA." *J. Esp. Zoology* **288**, 242-283, 2000.

[2] Alon, U., Surette, M.G., Barkai, N., Leibler, S. "Robustness in bacterial chemotaxis." *Nature* 397, 168–171, 1999.

[3] Blahut, R.E. *Theory and practice of error-control codes,* Addison-Wesley, 1983.

[4] Carlson, J.M. and Doyle, J. "Highly Optimized Tolerance: A Mechanism for Power Laws in Designed Systems," *Phys Rev E* 60, 1412 (1999).

[5] Carpenter, S. "Robustness in Ecosystems," SFI Robustness Site http://discuss.santafe.edu/robustness RS-2001-020.

[6] Dawkins, R. Selfish Gene, Oxford, 1976.

[7] Doyle, J. and Carlson, J. "Robustness and complexity," to appear in Robust Design, E. Jen (ed).

[8] Dierker, U., Hubler, A., and Dueweke, M. "Self-assembling Electrical Connections Based on the Principle of Minimum Resistance." *Physics Review* E 54, 496, 1996.

[9] Diacu, F. and Holmes, P. Celestial Encounters, Princeton University Press, 1996.

[10] DiMaggio, Paul J., and Walter W. Powell. 1983. “The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields.” *American Sociological Review* 48: 147-160.

[11] Hartman, J.L., Garvik, B., Hartwell, L. “Principles for the buffering of genetic variation,” *Science* 291, 1001–1004, 2001.

[12] Hubler, A. and Tracy, C. “Robust unstable systems” (private communication).

[13] Jen, E. “Robustness in social processes” (in preparation).

[14] Kauffman, S. Investigations, Oxford University Press, 2000.

[15] Kirschner, M. and Gerhart, J. “Evolvability,” PNAS 95: 8420–8427, 1998.

[16] Kearns, M. “Network security” (in preparation).

[17] Krakauer, D.C. Genetic redundancy. In Press: Encyclopedia of Evolution. Ed Mark Pagel. Oxford University Press. (In Press 2002)

[18] Krakauer, D.C. “Principles of robust design,” to appear in Robust Design, E. Jen (ed).

[19] Levin, S. Fragile Dominion: Complexity and the Commons, Perseus Books, 1999.

[20] Levin, S. “Robustness and Structure in Ecosystems,” SFI Robustness Site http://discuss.santafe.edu/robustness RS-2001-11.

[21] Levin, S. “Resilience in natural and socioeconomic systems,” *Environment and Developmental Economics*, 3:225-236, 1998.

[22] Low, B.S., Ostrom, E., Simon, C.P., Wilson, J. “Redundancy and diversity in governing and managing common-pool resources,” Eighth Conference of the International Association for the Study of Common Property, 2000.

[23] Marder, E. and Calabrese, R.L. “Principles of rhythmic motor pattern generation,” *Physiological Reviews* **76**: 687–717 (1996).

[24] Marder, E. “Robustness in neuronal systems: The balance between homeostasis, plasticity, and modulation,” to appear in Robust Design, E. Jen (ed).

[25] Padgett, J. and Ansell, C. “Robust action and the rise of the Medici, 1400-1434,” *American J. Sociology* **98**, 6: 1259–1319, 1993.

[26] Schneider, F. “Robustness in computing systems,” SFI Robustness Site http://discuss.santafe.edu/robustness RS-2001-036.

[27] Stadler, B., Stadler, P., Wagner, G., Fontana, W. “The topology of the possible, “ to appear in *J. Exp. Zoology.*

[28] Stark, D. "Heterarchy: Distributing Authorithy and Organizing Diversity." in *The Biology of Business: Decoding the Natural Laws of Enterprise.* J.H. Clippinger III (Ed.). Jossey-Bass Publishers, pp.153-179, 1999.

[29] van Nimwegen, E., Crutchfield, J.P. and Huynen, M. "Neutral evolution of mutational robustness." *Proc. Natl. Acad. Sci. USA*, 96: 9716–9720, 1999.

[30] Voigt, C. and Arnold, F. "Utilizing robustness to accelerate protein engineering," to appear in Robust Design, E. Jen (ed).

[31] von Neumann, J. "Probabilistic logics and the synthesis of reliable organisms from unreliable components," in *Automata Studies*, C. Shannon and J. McCarthy (eds.), Princeton University Press, 1956.

[32] Wagner, A. "Mutational robustness in genetic networks of yeast," *Nature Genetics* **24**, 355-361, 2000.

[33] Wagner, G., Mezey, J. and Calabretta, R. "Natural selection and the origin of modules," in W. Callabaut and D. Rasskin-Gutman (eds.) Understanding the development and evolution of complex natural systems, MIT Press, 2001.

[34] Webb, C. and Levin, S. "Cross-system perspectives on the ecology and evolution ecosystems," in Robust Design, E. Jen (ed.) to appear.

[35] Wiggins, S. Introduction to Applied Nonlinear Dynamical Systems and Chaos, Springer, 1990.

[36] Zweig, G. Private communication.