The "Agency" Model of Transactions: Toward an Understanding of Children's Theory of Control.

Edith K. Ackermann

Edith@media.mit.edu

Reference: Ackermann, E. In *Psychologie Génétique et Sciences Cognitives.* (Eds. Montangero, J. & Tryphon, A.). Fondation Archives Jean Piaget, pp. 63-73, 1991.[1]

Abstract

In this paper, I discuss how children come to disentangle *purpose* and *causation*—or psychological and physical descriptions—when explaining the behaviors of "simple-minded intelligent artifacts," as well as the behaviors of people and of things. I argue that very early on, children attempt to build a synthesis between these two kinds of explanations, and that they do so in a similar way when explaining the functioning of people, of objects, and of living and artificial creatures. The most striking characteristic of this synthesis is children's focus on "who" is impacting "whom" (and how), and "who" is impacted by "whom" (and how) in a given transaction. Both people and things are understood in terms of how they *control* each other's behavior, either through direct or mediated action (A *does* something to B, or *tells or signals* something to B, and B *acts or signals* back accordingly). Such = an "agency" model of transactions requires that the actors at play be animated. It forms the core of what will become a cybernetic view of the world. And like cybernetics, it provides a conceptual framework for explaining control and communication across the range of the living and the inanimate. Young children's tendency to *animate* objects constitutes a crucial step toward the construction of cybernetic theories. It is by no means a mere sign of cognitive immaturity.

---

[1] This chapter is based on an Address given at the 11th Annual Meeting of the Archives Jean Piaget on Genetic Psychology and Cognitive Sciences, University of Geneva, Switzerland, September 1990. Note: This paper also appears in: Constructionism. (Harel, I & Papert, S., Eds). Part 4: Cybernetics and Constructionism. Chap. 18, pp. 367-379. Norwood, New Jersey: Ablex Publishing Corporation, 1991

.

Introduction

Cybernetics was defined by Norbert Wiener as the study of "control and communication in animals and machines." (Wiener, 1952). First-order cybernetics provides a description of regulatory processes in terms of circularities, feedback, and retroactions. Second-order cybernetics describes the evolution of complex systems in terms of self-organizing principles, emerging behavior, and equilibration processes. Researchers in cybernetics explore these phenomena by building working artifacts.

In exploring cybernetic ideas with children, the purpose of the Epistemology and Learning Group is to develop and explore languages for describing the functioning of "simple-minded intelligent artifacts.". We want to compare languages used by children with languages developed by scientists during the early days of cybernetics. We want to contrast descriptions produced while children observe an artifact's behavior ("playing psychologist") with descriptions useful to their building and programming ("playing engineer"). Finally, we want to compare languages used to describe "intelligent artifacts" (objects that act in a people-like way) with languages used to describe living creatures and inanimate objects, or things. Such a project is no doubt ambitious and requires a group effort in its realization.

Many reasons have lead to our choice of this topic, not the least of which is technological. At M.I.T., the time has come when we can offer children a laboratory for building self-regulating devices. By "self-regulating device," I mean a mechanism able to read or *sense* certain features in its environment (such as light, sound, or obstacles), to measure or *evaluate* their value in relation to an internally fixed referent, and to *adjust its behavio* accordingly.

Our group has developed the building blocks out of which many kinds of "artificial creatures" can be built. The “creatures" are usually made out of LEGO bricks, gears, motors, and sensors, and controlled by a special version of the Logo programming language. Students in the Epistemology and Learning Group have created a variety of simple-minded sensori-motor artifacts, capable of all sorts of goal-seeking and regulatory behaviors, such as sensing light, avoiding obstacles, and responding to sound. More recently, Fred Martin and others have designed a set of "electronic bricks" that can be wired directly onto the creatures, allowing for their "autonomy" from the desktop computer. While Nira Farber and Fred Martin focus on the functioning of single "creatures" in interaction with their environment (first-order cybernetics), Mitchel Resnick is currently building a computer-based microworld for exploring the behavior of very large populations of such simple-minded creatures (second-order cybernetics).

The children we work with (mostly upper elementary and high-school students) are fluent enough in programming to design their own creatures (sometimes with the assistance of an adult). Under such circumstances, both experimenters and children find themselves in unusually rich terrain for exploring the functioning of artificial creatures. Not only can children *observe* the creatures' behavior as they are moving around in a natural or "social" environment (usually a box with obstacles, lights, sounds, as well as other creatures), but they can *influence* their behavior by acting upon their sensori imput. Even more important, they can *modify* their behavior *from within*, by reprogramming or rewiring them. In other words, children can switch roles from "playing engineer" to "playing psychologist"—from being the creatures' designer to being their external observer.

In this paper, I focus on a specific issue that attracted my attention as a developmental psychologist. I analyze how children gradually come to disentangle purpose and causation, or psychological and physical explanations, when describing the behaviors of simple-minded intelligent artifacts, as well as the behaviors of people and of things.

Of particular interest to me are questions such as: Under which circumstances are children likely to favor psychological or physical explanations? When and how do children distinguish between the behavior of animate and inanimate creatures? How do children gradually integrate causality and animacy when describing the behavior of "intelligent artifacts," which have the peculiarity, as mentioned before, of being an *object* yet acting in a *person-like (or animal-like)* way?

A survey of other researchers' work, together with my own preliminary results, indicates that very early on, children actually construct a "synthesis" of these two kinds of explanations to make sense of the functioning of—and transactions among—all people, things, and artifacts. The main characteristic of this synthesis is the description of a transaction in terms of "who" impacts "whom" (and how), and "who" is impacted by "whom" (and how). Causal chains are, in other words understood in terms of how elements control one another's behavior, either through direct or mediated action (A *does* something to B, versus A *tells or signals* something to B).

Such a quest for "agency" requires that children animate the actors in a transaction. The "agency" model, I claim, forms the core of what will become a cybernetic view of the world. And, like cybernetics, it provides a conceptual framework for explaining behavior across the range of the living and the inanimate. The ability to animate objects is a crucial step toward the construction of cybernetic theories. It is by no means a mere sign of cognitive immaturity.

Causality and Animacy among Cyberneticians

We know from the history of cybernetics (and later, of AI and cognitive science) that as soon as the first intelligent artifacts were built, many fundamental concepts such as "self," "purpose," "intentionality," and "free will," needed to be reconsidered. The reason for this questioning is the profoundly ambiguous nature of these early "intelligent" artifacts. They are physical objects, yet they behave in a somewhat human-like fashion. Their dual nature led both their builders (engineers) and their observers (psychologists) to engage in heated debates about the nature of their "intelligence." Could they be called intelligent (in the first place)? And if so, how could one best describe their intelligence?

The Macy Conferences provide the richest existing trace of the spirit in which these early discussions evolved (In H.Von Foerster, (Ed.) 1949-1952). In studying the Macy Conferences, we noticed time and again that engineers and psychologists quite consistently picked different sets of concepts for describing the functioning of self-regulating devices. More recent debates within the fields of AI and cognitive science seem to confirm the polarities already observed during the Macy Conferences (Boden, 1978; Dennett, 1987; Minsky, 1986). The story goes as follows:

Through actually building intelligent systems out of unintelligent parts, engineers make it clear that goal oriented-ness (or knowing how to reach a goal) by no means requires the system's awareness (its knowing that it knows how to...), or intentionality (its deciding to...), or free will (if willing, it could...). Engineers usually insist that no vital principle needs to transcend, or live independently, from a material substrate. To them, if mind might well emerge from matter, the building of a mind—out of matter—does not require the use of higher-order concepts such as purpose and intentionality. Such concepts are not operational or useful in the work of engineers.

In contrast, most psychologists (as well as educators) study behavior *as it becomes meaningful to—and controllable by—a subject*. And not surprisingly, since it is their job to help people use whatever reflecting capabilities they have (self-awareness) as a means of monitoring their own behavior. To the psychologist, the concept of "subject" is a useful construct, and so are the concepts of "purpose" and "intentionality." Their role is to capture at a glance what seems so specific to human cognition—namely, the ability to act upon its own activity as if it were a "subject" of inquiry, to evoke objects that are not physically present, to act (mentally or actually) upon objects (real or symbolic) , and to choose the most appropriate among a set of possibilities to achieve a goal. Even Piaget, whose interest as a structuralist was in studying "what people *know how to do*" rather than on "what they *think they know how to do* " (Piaget's own words) spent many years of his life trying to understand how babies gradually outgrow early reflex activity and become able to initiate, adjust, and inhibit action at their own will (Ackermann, 1990). In the enormous progress achieved during the first two years of their lives (the sensori-motor period), babies no doubt move from being entirely driven by many uncontrollable needs, to partially monitoring their drives through self-correcting activities (circular reactions) and later, through reflection (reflective abstraction). This transition from *reflex* to *reflective* activities, or from response to command, needs to somehow be captured by the cognitive psychologist.

Causality and Animacy in Children

Developmental psychologists generally assume that cognitive growth starts with an initial state of in-differentiation between a subject and her environment (or of a dualism between the self and the outside world), and that it moves progressively toward a state of greater differentiation. The subject is thus enabled to perceive and describe the world (things and others) as being separate from herself. If such a view is generally accepted, the following two questions nevertheless remain:

1) Is this process of differentiation is different when dealing with things and with people (or with inanimate and animate objects)?

2) Is this process is as smoothly incremental and as universal as developmental psychologists generally assume?

People's ability to progressively disentangle purpose and causality gains to be situated within this broader developmental framework. Research in this domain leads to a somewhat contradictory picture.

On one hand, psychologists like Michotte (1963), Leslie (1979-1984) Heider & Simmel (1944), and Steward (1982) suggest that one of the primary distinctions that people make in understanding and dealing with the world is between things they react to as social objects (people-like), and things that they react to as non-social (things-like.) Very roughly speaking, social objects—or people-like things—are assumed to be active agents, capable of initiating behavior from within and of directing, inhibiting, and modifying their behavior after an analysis of its effects or consequences. In other words, social objects are perceived as having some *controlling intelligence* that helps them to monitor their activity. In constrast, non-social objects—or things—are assumed to be passive receivers, or recipients of others' actions. The action of a thing is non-intentional and not initiated from within. Instead it is driven by external forces, which act upon it. *Actions happen to them, rather than them causing actions to happen.* This group of researchers claims that people recognize quickly, and at a young age, the difference between the functioning of social and non-social objects, and that people then regulate their behavior accordingly. These researchers consider that people usually describe social objects in psychological terms, and non-social objects in mechanical-causal terms.

On the other hand, we know from research by Piaget (1975), Carey (1985), DiSessa (1982, 1988), Inagaki & Hatano (1987), and Turkle (1984) that young children do not, most of the time, distinguish mind from matter, or psychological from physical reality. And consequently, many things, such as moving objects, appear to them as endowed with both material properties and with will. To use Piaget's words:

> In early stages of their development, children do not draw limits between themselves and the external world. This leads us to expect that children will regard as living and conscious a large number of objects which are, for us, inanimate (Piaget, 19XX, p. 77).

Actions such as movements of an inanimate object are often seen as a result of a conscious effort. And if an agent *knows how* to do something, young children assume that it also *knows that it knows how* to do something.

Causality and Animacy: Existing Research

In the face of such apparently contradictory results, it is useful to look more closely into the existing body of research, in an attempt to develop an explanation for the apparent disparities. Such a preliminary survey is a necessary prerequisite to analyzing our own data. It has actually provided me with a model, or rather with a lens, through which I could make sense of the children's attempts to disentangle the concepts of animacy and causality.

> "There is a celebrated monograph, little known outside academic psychology, written a generation ago by the Belgian student of perception Baron Michotte. By cinematic means (shapes moving on a screen) he demonstrated that when objects move with respect to one another, within highly limited constraints, we see causality" (Bruner, 1986, p.17).

Let me briefly describe Michotte's experiment on the perception of causality:

• In a first situation (1), a shape A moves toward a shape B, makes contact with B (hits it), stops, and then B moves away. All subjects immediately describe shape A as launching shape B.

• In a second situation (2), the action begins in the same way. Yet after A contacts B, both shapes continue in the same direction. In this case, shape A is described as dragging the other.

• If, as in the third situation (3), the time span increases between the moment when A hits B and the departure of B, people tend to split the episode into two parts: B is no longer seen as responding to the action of A, but is described as initiating its own movement.

Michotte's experiment was initially designed for adults. It has since been replicated with young children (Piaget & Lambercier, 1958) and with infants (Leslie, 1979, 1984). Leslie was interested in the origins of such perception of causality in very early childhood. His experimental procedure consisted of measuring signs of surprise in six-months-old infant. Leslie showed infants a sequence of cinematic episodes that we, as adults, see as being caused. He then interspersed one non-causal episode (sequences showing action at distance), and the infants, he claims, showed startled surprise.

Yet another variation of Michotte's experiments, imagined by Fritz Heider and Marianne Simmel (1944), was to use a "bare" animated film to demonstrate the immediacy of "perceived intention." They presented adults with a scenario made out of a small moving triangle, a large moving circle, and a box-like empty rectangle. All the shapes moved in and around the empty box. Observers unanimously described the movements as a series of animate moves with clear social intent, like chase, fight, protection, attack, etc. According to Bruner, the scene was usually perceived by adults as two friends being purchased by a large bully-type of creature, who, upon being thwarted, broke up the house in which he had tried to find them. All the scenarios produced instances of social causation.

Finally, Judith Steward brought Michotte's experiment a step further, by showing that very specific properties of movement systematically trigger intentional rather than causal interpretations. In an experiment called "Object motion and perception of animacy," Steward presented adults with 35 displays programmed on an Apple II computer. The displays showed moving dots in a variety of situations.

Either the source of a dot's movement is unknown (a dot just starts) or it is triggered by something else. In some cases, the dots move together in a straight line. In other cases, they don't. One dot might overtake another through sudden acceleration, or it might decelerate and stop. Dots might change the course of their trajectory, and make detours around obstacles. Steward has varied the spatiotemporal constraints of "blob" movements in such a way as to produce apparent "animacy" or "causality." Her question to the subject was very simple: "I have done a series of simulations of movements of objects and creatures. I will show them to you one by one, and you tell me if you think they are objects, creatures, or whether you cannot tell."

Steward shows that subjects plainly see goal-seeking persistence when dots initiate behavior, when they overcome obstacles, and when they accelerate and change the course of their trajectories. Inanimate objects, in contrast, are supposed to conserve their heading, their direction, and their speed. If they do not conserve their speed, inanimate objects can only decelerate (or die out).

Note that with the exception of Leslie, this first group of researchers essentially studied adults. Their findings are nonetheless widely used to show that people recognize quickly, and at a young age, the difference between the functioning of social and non-social objects, and that people then regulate their behavior accordingly (Bruner, 1986).

Research by Piaget, Carey, DiSessa, and Iganaki & Hatano, on the other hand, shows that purpose and causation are not always as clearly distinguished as suggested by the first group of researchers. Young children, as well as adults, tend to personify objects and, at times, to objectify persons. In Piaget's study on animism, young children quite generally express the idea that objects, such as a stone or a piece of wood, cannot be hit or broken without feeling the impact. In a similar way, a watch cannot give the time without knowing that it does so; and a boat cannot carry people without making an effort and, of course, feeling that it is making an effort.

A more careful reading of children's protocols however shows that their analysis of the circumstances under which an object is likely to feel an impact is extremely subtle. My favorite example comes from Ali, age eight and a half. The experimenter asks Ali whether he thinks that a cloud could feel things —like, for example, a sting. The child answers that, of course, a cloud cannot not feel a sting. How could it since it is only air? Later in the interview, however, Ali comes back and mentions that clouds actually do feel wind and warmth, since it is both wind and warmth that drives them. Another child, around the same age, thinks that a wall cannot feel anything ("because it is strong!"). It is only when someone hits it really hard and knocks it down that it feels something ("because THAT would break it"). The point here is that if children may well be animistic, they are so in an extremely selective fashion. The interaction between elements needs to make sense, and moreover, the personification of an element needs to respect and possibly explain physical laws or psychological principles (Hatano & Iganaki, In preparation).

What does this apparently conflicting body of research tell us about children's ability to disentangle purpose and causation?

First, it reminds us that intentions do indeed constitute a fundamental and quite primitive category with which personal experience and transactions between objects (alive or not) are organized. Young children tend to view physically caused events as being psychologically intended. This does not imply that children simply project human-like properties to all physical or biological "agents" indifferently. By "agent" I mean an element (alive or not) that is currently acting upon another element during a transaction. Children introduce increasingly subtle distinctions between specific situations. They attribute more or less control (will) depending on whether they see an "agent" as being omnipotent—and omniscient— and also, on whether they perceive the limits of their own ability to impose their will over people and things (Iganaki & Hatano, 1987).

The examined body of research, moreover, helps us identify specific properties of an "agent's" behavior which trigger purposive rather than causal explanations (Michotte, 1963; Steward, 1982). If a "blob" initiates a movement by itself, if it "knows" how to change direction or speed along its way, and/or if it "insists" on a goal-seeking motion, both children and adults tend to "animate" it, or to describe it in psychological terms. They give it a soul (anima),or a brain, to account for the self-regulatory aspect of its behavior. On the other hand, if a "blob" initiates movement and changes direction or speed under the impact of another "blob," both children and adults tend to "objectify" it. One could say that such a "blob" has no inner locus of control, and thus no "anima."

However, our own pilot studies—as well as a closer look into children's responses to Michotte's and Heider & Simmel's experiments—indicate that causality and animacy are not always as clearly differentiated (as suggested by Michotte and others). In a recent doctoral thesis, Thommen (1990) studied children's perception of "intentional causality" (her words), when presented with Heider and Simmel's animated film. Her analysis shows that children, between age three to twelve, progressively refine their distinctions between purpose and causation when describing the scenario (See also Gilliéron & Thommen, 1987)

Moreover, in analyzing the exact terms used by young subjects in Michotte's experiment (Piaget & Lambercier, 1958), and in asking them to be explicit about the prerequisits and consequences of each "blob"'s activity, we see that children certainly "mingle" intentions and causes, yet in a very consistent way. They animate or dis-animate agents as they focus on particular aspects of their activity.

The Agency Model of Transactions: A Step toward a Theory of Control

Children's explanations of mechanical-causal realities, as well as of biological and psychological realities, are remarkably consistent across the range. They are best captured by what I call an "agency model of transactions." Such a model becomes progressively refined during ontogeny. It varies across situations, and evolves as a transaction occurs. The story goes as follows: Imagine an "agent" A (alive or not alive) acting upon a "recipient "B, either by directly *doing* something to B (impacting it) or by *telling* it something (giving it a command). If the child does not know—or pay attention to—what has triggered "A's" behavior in the first place, then A tends to be seen as an active agent initiating behavior "at its own will". And usually, as A becomes animated, B gets objectified. Now as soon as you ask children to think about the consequences of A's behavior, they turn their attention to the recipient B, and begin to envisage its ability to respond. My claim is that as soon as this shift of attention takes place, young children actually "animate" B. Instead of being a simple recipient, B thus turns into an "agent." The child focuses on *its* feelings and *its* ability to act, or react, and by doing so, dismisses B's initial object-like passivity. Once focused, B gains life and autonomy. It can now feel an impact, bounce off, hit other elements, and hit back.

Note that it is only insofar as B is able to react to A that A's action can be described as a signal. In effect, how could a signal "flow" from A to B, if A's action (physical impact or command) did not affect B's behavior somehow? If B "acts forward" (impacting a third element C), we—researchers—usually describe the event as a linear "causal" chain. If instead B "acts back" on A, modifying its behavior, we think of the event as circular causality, or feedback. In both cases, the actors at play (alive or not alive, agents or recipients) can be said to *control each other's behavior*. They do so either via a physical impact (A *does* something to B) or via instructions (A *tells* something to B). In the first case, we—researchers—tend to describe the transaction as a case of physical or mechanical causality. In the second case, we talk about information or signal processing. Unlike adult researchers, young children do not seem to distinguish between mechanical causality and signal processing. Instead, they consistently animate agents under scrutiny. And since recipients (who were previously objectified) can themselves become "alive", children are able to view a transaction from a different perspective. They learn to describe how different elements (actors) affect or *control* one another's elements' behavior. Whether these actors are alive or not alive, or whether they act—and affect a recipient's behavior—by actually "pushing" it (physical impact) or by sending it a message (information) seems less relevant. The "agency" model, on the other hand, facilitates the distinction between linear and circular causality. Such a distinction becomes possible as soon as the child understands that a given agent (B) could both *act forward* by impacting a third element (C), and *act backward* by reacting to a first element (A).

Many questions can be asked based upon the agency model of transactions. To name but a few: What triggered an agent's behavior? Does an agent recognize the consequence(s) of its own activity? If it does, can it modify its behavior accordingly? can a a recipient be sensitive to an agent's impact? How and why will it respond? In answering these questions, children refine their descriptions of a transaction, while at the same time, unveiling how they articulate purpose and causation. Note that very often children set these questions for themselves without an adult's intervention.

The objective of my current research on children's development of cybernetics explanations is to probe the explanatory power of the agency model in a variety of settings. I use "agencies" as a framework for designing experiments, for generating questions and, above all, as a lens for interpreting children's spontaneous descriptions. My long detour into children's theories of human, and non-human, transactions was a prerequisit to studying their developing theories of minds (their own and other creatures' minds). When describing the functioning of simple-minded artificial creatures (interacting with their environments, or among themselves) children do, indeed, think in terms of agencies (of "who" is controlling "whom" or "what"). And agencies, I suggest, are an essential step toward a cybernetic view of the world.

Acknowledgments

I thank Idit Harel, Mitchel Resnick, and Carol Strohecker for their help in writing this article. I am grateful to Bärbel Inhelder, Guy Cellérier, Andy DiSessa, Heinz von Foerster, Christiane Gillieron, Ernst von Glasersfeld, Seymour Papert, Jean Piaget, Marvin Minsky, Evelyne Thommen, and Mike Travers for their inspiration. Above all, I thank my colleagues at the Epistemology and Learning Group, my students, and my former colleagues from the F.P.S.E., in Geneva. I especially wish to thank the children whose creative and generous minds are a main motivator in all my work.

References

Ackermann, E. (1990). Circular reactions and sensori-motor intelligence: When Piaget's theory meets cognitive models. In *Archives de Psychologie*. Vol 58, 65-78.

Boden, M. (1978). *Purposive explanation in psychology*. England: The Harvester Press.

Bruner, J. (1986). *Actual minds, possible worlds*. Cambridge: Harvard University Press.

Carey, S. (1985). *Conceptual change in childhood*. Cambridge: MIT Press.

Dennett, D.C. (1987). *The intentional stance*. Cambridge: MIT Press.

DiSessa, A. (1982). Unlearning Aristotelian Physics: A study of knowledge-based learning. In *Cognitive Science*,. N. 6, 37-75.

DiSessa, A. (1988). Knowledge in pieces. In Forman & Pufall (Eds). *Constructivism in the computer age*. New Jersey: Lawrence & Erlbaum.

Epistemoology an Learning Group. (1990). Harel, I. (Ed) *Constuctionist Learning*. A Media Laboratory Publication. Cambridge: MIT.

Gilliéron, C. & Thommen, E. (1987). La perception de la causalité sociale chez l'enfant de 3 à 9 ans. In *Psychologie et education*. N.3/4, 158-161.

Hatano, G.& Inagaki, K. (In Preparation). Everyday biology and school biology: How do they interact?. In Paris & Stevenson (Eds). *Cultural perspectives on litteracy, cognition and schooling*.

Heider, F. & Simmel, M. (1944). An experimental study of apparent behavior. In *American Journal of Psychology*, 57, 243-259.

Inagaki, K. & Hatano, G. (1987) Young children's spontaneous personification as analogy. In *Child development*. N. 58.

Leslie, A.M. (1979). *The representation of perceived causal connection*. Doctoral Thesis, Department of Experimental Psychology,. University of Oxford.

Leslie, A.M. (1984). Spatiotemporal contiguity and the perception of causality in infants. In *Perception,* 13, 287-305.

Michotte, A. (1963). *The perception of causality*. New York: Basics Books.

Papert, S.(1980). *Mindstorms: Children computers and powerful ideas*. New York: Basic Books.

Piaget, J. (1975) . *The child's conception of the world*. New Jersey: Littlefiel, Adams. (Note: The discussion on animism is found in p.194-206).

Piaget, J. & Lambercier, M. (1958). *La causalité visuelle perceptive chez l'enfant et chez l'adulte*. In Archives de Psychologie, 36, 77-201.

Steward, J. (1982). *Perception of animacy*. Doctoral Thesis. University of Pennsylvania.

Thommen, E.(1990). *La perception de la causalité intentionnelle chez l'enfant de trois à douze ans*. Doctoral Thesis. University of Geneva: F.P.S.E.

Turkle, S. (1984). *The second self*. New York: Simon & Schuster.

Von Foerster, H. , Mead, M. & Teuber L. (Ed.) (1949-53). *Cybernetics: Circular, causal, and feedback mechanisms in biological and social systems*. Transcripts of the 6th, 7th, 8th and 9th Josiah Macy Conference. Caldwell, NJ: Progress Associates, Inc.

Wiener, N. (1952). *Cybernetics, or control and communication in the animal and the machine*. New York: John Wiley & Sons.

Wiener, N. (1954). *The human use of human beings: Cybernetics and Society*. New York: Garden City.