Journal of the Association for Information Systems (2018) **19**(10)**,** 929-952
doi: 10.17705/1jais.00513

RESEARCH PAPER

ISSN 1536-9323

When Humans Using the IT Artifact Becomes IT Using the Human Artifact

Dionysios S. Demetis[1], Allen S. Lee[2]

[1]Centre for Systems Studies, Hull University Business School, U.K., d.demetis@hull.ac.uk
[2]Virginia Commonwealth University, U.S.A., allenslee@alum.mit.edu

Abstract

Following Demetis & Lee (2016) who showed how systems theorizing can be conducted on the basis of a few systems principles, in this conceptual paper, we apply these principles to theorize about the systemic character of technology and investigate the role reversal in the relationship between humans and technology. By applying systems-theoretical requirements outlined by Demetis & Lee, we examine conditions for the systemic character of technology and, based on our theoretical discussion, we argue that humans can now be considered artifacts shaped and used by the (system of) technology rather than vice versa. We argue that the role reversal has considerable implications for the field of information systems that has thus far focused only on the use of the IT artifact by humans. We illustrate these ideas with empirical material from a well-known case from the financial markets: the collapse ("Flash Crash") of the Dow Jones Industrial Average.

Keywords: IT artifact, Human Artifact, Systems Theory, Systemic Technology

Jason Thatcher was the accepting senior editor. This research article was submitted on June 7, 2017 and went through two revisions.

1 Introduction

The field of information systems (IS) rests largely on examining the contextual use of technology within social (sub)systems and organizations. In such a context, the relationship between the social and the technical has always been of special interest to IS researchers. Ultimately, this interest is applied to the interactions between humans and information technology, and at the center of attention, one can often find the concept of the IT artifact. From considering the IT artifact as an ensemble of hardware and software (March & Smith, 1995) to bundles of material and cultural properties that are recognizable and emerge from ongoing socioeconomic practices (Orlikowski & Iacono, 2001), or even to sociotechnical assemblages (Silver & Markus, 2013), one thing is clear: the concept of the IT artifact has mutated substantially over the years in how it has been depicted by the IS community. In fact, the ontological dimensions upon which the "IT artifact" has come to be considered have shifted so much that Steven Alter's suggestion was to "retire" that concept altogether from the lexicon of IS scholarly debate, as it had outlived its usefulness (Alter, 2015). While this is not an essay about the concept of the IT artifact per se, we do make the argument that the nebulous character of that concept is due to a much larger (though subtle) phenomenon at play: the transition of technology from *artifact* to *system.* In fact, the contextual richness that has been added to the concept of the "IT artifact" after its first use by March and Smith (1995), can be reinterpreted as a recognition of such a transition. Inspired by the posthumanist tradition that reflects on the boundaries between humans and technology, we take a different approach. We develop a *systems theoretical* description of the transition from artifact to *system* and argue that people are becoming agents of (the *system* of) technology.

Over time, technology has penetrated society to such a degree that even basic functions now seem almost inconceivable without technology. Indeed, this level of societal dependence on technology has become so deep that—in a large number of fields—there are now no manual fallback plans in cases of technological failure. By and large, even when technology fails, we tend to rely on more technology for rectifying the problems of technological use. Also, the rising trend of technologized decision-making that has taken certain fields by storm is even more alarming. In the foreign exchange markets, for example, 85 percent of all trading is conducted by algorithms alone, i.e., without any human intervention; this led the scholars investigating the phenomenon to call it the "Rise of the Machines" (Chaboud, Chiquoine, Hjalmarsson, & Vega, 2009). In the UK, the "ultra-high-speed version of algorithmic trading, high frequency trading, is estimated to account for over 77% of transactions in the UK market" (Sornette & Becke, 2011, p. 5).

A skeptic of our position, who might seek to argue against the trend of technologized decision-making, might pose this question: *Is it not the case that the designers of algorithms are humans?* And if so, then couldn't someone consider the role of algorithms (and of technology at large) as an extended application of human decisions?

Our challenge is to convince the reader of the contrary. For this purpose, we address this issue through a few intertwined questions: How does technology subvert and subdue human decisions? What conditions can be identified (with the help of systems theorizing) for this new role that technology has assumed, and how does this constitute the emergence of a *system of technology*? Even more crucially, how is it that humans assume the role of "artifacts" being shaped and used by technology in this—seemingly counterintuitive—role reversal? For the latter condition, we reserve the term "artificial humans" or "human artifacts"—in contrast to Simon's (1996) "artificial things" (i.e., artifacts).

As we will see through our example of the Flash Crash regarding the Dow Jones Industrial Average, the role of technology leads us to consider a seemingly radical idea at first. This idea, we argue, is an accurate reflection of how technology shapes social systems and subjects humans to forces that cast them "out to the environment" (in a systems theoretical sense, whereby the environment is what lies outside the boundary of a system), outside of what has become technologized decision-making; instead of the IT artifact being shaped and used by humans, humans can actually be considered as "artifacts" being shaped and used by machines. By casting humans out to the environment, technology does not make humans redundant altogether. In fact, the human artifact, as we incorporate it into our theorizing, is a role that humans assimilate themselves into. So an *environment* surrounding a *system* of technology is not, from a systems theoretical perspective, an ontological isolation chamber. A boundary between a system and its environment creates a distinction that is helpful analytically, but relations across the boundary are not eliminated. As Luhmann indicates:

> *The concept of the environment should not be misunderstood as a kind of residual category. Instead, relationship to the environment is constitutive in system formation. It does not have merely "accidental" significance, in comparison with the "essence" of the system. Nor is the environment significant only for "preserving" the system, for supplying energy and information . . .the environment is, rather, a presupposition for the system's identity, because identity is possible only by difference. . . . Everything that happens belongs to a system (or to many systems) and always at the same time to the environment of other systems. (*Luhmann, 1995, pp. 176–177)

Thus, even with technology dominating decision-making in certain fields (e.g., finance), human/technology relations continue to occur but we argue that while human agency is reduced, the reconfiguration of the relations between humans and technology is guided largely by the *emergence of a system of technology*. This is qualitatively different from a quest to clarify the concept of the IT artifact per se and inspired by a theoretical tradition (Bateson, 1972) that accepts "that the point from which all further investigations in systems theory must begin is therefore not identity but difference" (Luhmann, 1995, p. 177). In this view, and in domains where human agency is becoming subordinate to automated executions, it is humans that must react to *technological stimuli* rather than technology that must react to human stimuli. Furthermore, the technological stimuli are emergent and not predesigned (or preprogrammed) in any way. This also assumes that while the controllability of technology can be achieved at a microscale (where one could assert that the link between *designers* and (control of) *artifacts* is strict), at a macroscale, technology exhibits emergent nonlinear phenomena that render controllability infeasible. Ultimately, this transition from controllability at the microlevel (the domain of computer science) to emergent and systemic nonlinearity at the macrolevel showcases the pressing need for the field of IS to explore the much larger social, economic, cultural, and organizational shifts that reduce human agency and result in what we call a role reversal between humans and technology.

Stripped of causality and linearity at the macrolevel, as well as devoid of controllability, technology emerges as a nondeterministic system of interference that shapes human behavior. In turn, humans *react* to the nondeterministic emergent stimuli that a system of technology spawns. Thus, our description demands a systemic role for technology, with humans increasingly finding themselves in the environment of that system with which they remain coupled—indeed, in "loose couplings" that often reduce humans to artifacts themselves.

The next section of this paper presents some related work and developments in the concept of the artifact that aim at exploring the significance of the broader trajectory of the IT artifact and the emergence of a system of technology. The third section of the paper discusses a few examples in brief while the fourth section provides a review of a selection of the general requirements proposed by Demetis & Lee (2016) for systems theorizing, for the purpose of theorizing about *technology itself as a system* in the systems theoretical sense. The fourth section also serves to pose reflective questions about the deep interference of technology in society and highlight elements that ought to be considered for technology in this context. The fifth section presents the case of the Flash Crash in the Dow Jones Industrial Average index and reviews the key characteristics of that case by considering how technology shaped the minicrisis through automated execution strategies. In the sixth section of the essay, we engage in a brief discussion on the consequences of the basic systemic principles of technology. Finally, in the seventh section of the paper, we offer our thoughts on a research agenda for the future.

2 Developments in the Concept of the "Artifact"

In this section, we seek to review some basic principles around the concept of the artifact but also explore what the broader trajectory of IT artifact variants implies. While it is beyond the scope of this paper to delineate all the different variations of the concept of the IT artifact, we would like to clarify that we are more interested in what the existence of all these different variants signifies. As we will argue, the IS community has thus far taken the approach of piling one "IT artifact" concept on top of another while overlooking what these variations in concepts signify. But while the IS field has been, perhaps, preoccupied with "micromanaging" the description of the concept of the IT artifact, the field has failed to consider the possibility of humans assuming the role of artifacts that are "cast out to the environment" of a system of technology. For a reflection on the evolution of the concept of the IT artifact and an exposition of the "artificial human" (i.e., the human artifact) following our case study, we turn to a handful of cardinal points on the nature of the artifact.

In their work going back to the origins of the concept of the artifact, Lee, Thomas, and Baskerville (2015) refer to Herbert Simon's work and, more specifically, to the indicia that demarcate the distinction between the natural and the artificial. We regard these indicia to be foundational in our argument about whether the IT or the human can be considered as an artifact. According to Simon's treatise (1996, p. 5), as quoted by Lee et al., these are portrayed as follows:

1. Artificial things ["artifacts"] are synthesized (though not always or usually with full forethought) by human beings.
2. Artificial things may imitate appearances in natural things while lacking, in one or many respects, the reality of the latter.
3. Artificial things can be characterized in terms of functions, goals, adaptation.
4. Artificial things are often discussed, particularly when they are being designed, in terms of imperatives as well as descriptives.

We will return to Simon's indicia when we later explicate our concept of the human artifact, in contrast to the IT artifact.

In the context of a transition from the general "artifact" concept to the "IT artifact," March and Smith (1995) take a design science approach informed by Simon's work and view IT artifacts as "instantiated in specific products . . . intended to perform certain tasks" (p. 253); IT artifacts embody specific characteristics and can be broadly understood as consisting of hardware and software. In addition, Simon (1996) clearly conceives of the computer as an artifact:

> *The computer is a member of an important family of artifacts called symbol systems, or more explicitly, physical symbol systems. . . symbol systems are almost the quintessential artifacts, for adaptivity to an environment is their whole raison d'être. They are goal-seeking information processing systems, usually enlisted in the service of the larger systems in which they are incorporated.* (p. 21–22)

Of course, despite the existence of some systemic characteristics in Simon's work in the context of the quotation above (e.g., system/environment), it was generally understood that computers as artifacts adapted to a very particular type of "environment"—one both created and guided by human decisions. Such artifacts are "human-created artifacts that have value insofar as they address this task" (March & Smith, 1995, p. 258). But as we shall argue in our discussion, and support through our case, this is gradually being reversed with IT artifacts developing into *systems* that largely *create the environment* to which they themselves react. In what can be recognized as a recursive process (Luhmann, 2002), IT artifacts subdue and confine human decisions and reduce their function in a way that—in turn—supports the IT artifacts themselves. This recursive condition gives rise to the additional role of the human artifact: a human of reduced agency that is both shaped by and reacts to technological stimuli that are emergent and nondeterministic.

March and Smith's reflections in the context of design science were followed in a more general context by Orlikowski and Iacono (2001), who emphasized that there is a significant undertheorizing of the concept of the IT artifact and saw the latter as superficially represented in IS research. The impact of Orlikowski's and Iacono's work is evident—their article was followed by the publication of "360 refereed articles" within a timeframe of almost 14 years (Lee et al., 2015, p. 6); however, this "surge of interest in technological artifacts was further accelerated" (Lee et al., 2015, p. 6) also by Hevner, March, Park, and Ram (2004). The shift that occurred in the reconceptualization of the IT artifact created another vision that emphasized the packaging of material and cultural properties in some recognizable form of hardware and/or software (Orlikowski & Iacono, 2001, p. 121), thereby highlighting the contextual significance within which IT artifacts are embedded. While it is beyond the scope of this paper to revisit the multiplicity of different conceptualizations of the IT artifact, we must nevertheless stress the ontological significance of such a multiplicity. Whether one invokes what Orlikowski and Iacono name a tool, proxy, ensemble, computational or nominal view of technology, one is essentially ignoring the recursive nature of how the artifact affects itself; this level of how technology interferes with itself (a self-referential expression) is hard for humans to contemplate (and even accept), as it leads to the counterintuitive proposition that humans are now the artifacts being shaped, a position that implies humanity is unwittingly adapting or being adapted on the basis of *unintended* computer-generated consequences. As we shall see, in some cases, future configurations of human behavior are even designed out of such unintended computer-based stimuli. This differs radically from other perspectives that conceptualize IT within a specific context and a structure embedded in that context (Benbasat & Zmud, 2003). While there is indeed a sociotechnical assemblage of IT artifacts (Silver & Markus, 2013) in what can be conceived of as a "sociotechnical artifact," we posit that technology is overtaking not only human decisions and the context of their embeddedness, but also entire subsystems of society (such as, as we will later show, the financial system).

The avalanche of IT artifact variants that have sprung from the work of Orlikowski and Iacono is—if considered as a unity and viewed from a metalevel—enough of a reason to make one doubt their usefulness. In turn, this variety has been interpreted as confusion and has even led scholars like Steven Alter to "retire" the concept of the IT artifact altogether. While we would not argue that the concept must be made redundant altogether, we would agree with Alter's insightful concerns and ask: what does the coexistence of so many models of IT artifacts tell us? How can we rethink the identity of the IT artifact through contingent differences?

While Simon's (1996, p. 3) assertion that "as our aims change, so too do our artifacts and vice versa" remains valid, the IS field has thus far ignored the "vice versa." What is this vice versa expression exactly? As our artifacts change, our human aims change too. While we cannot attribute to Simon an inference on a role reversal as we see it with the concept of the human artifact, we would add that our behaviors, intentionalities, and designs, change accordingly; they find themselves under the influence of nondeterministic effects of complex technologies. Thus, what we have is an inversion of influence between the IT artifact and humans.

One of the main reasons why the evolution of the IT artifact is hard to explore is because it is only one side to a contingent distinction (IT artifact/Human artifact); like any distinction, this distinction too is neither objective nor fixed but contingent upon an observer that decides how the distinction is to be explored. Any observer's choice can then be subjected to second-order observations (Luhmann, 1995); in simple terms this means that different observers are never looking at the same thing when observing. This is exponentially more challenging when we consider the ensemble of complex technologies that are supporting, as an example, the totality of the financial system. Can any one observer delineate how such an "artifact" even functions? To what degree has technology acquired a life of its own, so to speak? In such a context—and put simply—there can never be resolution or crystallization of the concept of the "IT artifact" because it is only one part of a contingent dynamic and whatever slice scholars explore, the whole remains unobservable (Angell & Demetis, 2010; Luhmann, 2002): we are looking at a continuous creation/recreation of technological blind spots, and the way this affects society is far from clear. While the recognizable contextual richness that has been added to the IT artifact by Silver and Markus (2013) alerts us to the importance of the context within which artifacts are embedded, the continuous metamorphosis of the IT artifact only serves to illustrate how different scholars have been building theoretical and conceptual constructs upon moving sand. While this is frustrating for the IS field that has weaved the quest for a conceptually lucid IT artifact into the very existence of IS (and the resolution of its identity crisis), we argue that this quest is misplaced. We should embrace this condition, instead of considering it as an identity crisis. As Foerster (1971, p. 1) reminds us: "Hard sciences are successful because they deal with the soft problems; the soft sciences are struggling because they deal with the hard problems." Ultimately, information systems, as a soft science, has taken the concept of the IT artifact as the field's dependent variable, succumbing to the Popperian delusion of objectivity in relation to empirical domains:

> *The empirical basis of objective science has thus nothing "absolute" about it. Science does not rest upon solid bedrock. The bold structure of its theories rises, as it were, above a swamp. It is like a building erected on piles. The piles are driven down from above into the swamp, but not down into any natural or "given" base; and if we stop driving the piles deeper it is not because we have reached firm ground. We simply stop when we are satisfied that the piles are firm enough to carry the structure, at least for the time being* (Popper, 2002, p. 94)

Applying Popper's ideas here to the concept of the IT artifact, we can say that we must deny any pretense to an ultimate objective foundation for such a concept. This applies also to the concept of the human artifact, discussed later on in a rudimentary form. In the transitions explored in the concept of the IT artifact, we (as an IS community) have merely been piling one concept on top of another, hoping that at some point we will hit Popper's "solid bedrock." Popper reminds us that even within the natural sciences, there is no solid bedrock. The reason that the process of accumulating artifact variants hasn't stopped is because the IS community is not satisfied that the piles are firm enough to carry the structure, all the while losing sight of how the swamp has been moving. The disruptive uncertainties that emerge from networked interactions and complex computations, raise their heads above any intertwined dependencies of sociotechnical and/or sociomaterial assemblages. Thus, throughout this process, the IS field has been driving the piles of "IT artifacts" above a swamp of transitory technological changes. What we are asking, in the context of this metaphor, is "how is the swamp" moving (in the Popperian sense) and what does that mean? This, we believe, is a more pertinent question with more critical theoretical consequences. By embracing this condition, we may explore the possibility of shifting the focus from IT as the artifact that humans are designing and shaping for human purposes, to humans as the artifact that IT, or a *system* of technology, is designing and shaping for the purposes of the *system* of technology. We recognize this systemic nature of technology in the systems theoretical sense and the ensuing role reversal as a unique type of subordination of human decision-making. We (as humans) now live in a *society of technology*, and in an increasing number of domains (such as finance), we assume the role of artifacts ourselves.

3 Examples of Technologized Decision-Making

Indeed, one can find a wealth of examples where technological autonomy has developed into a system that takes over important decisions—and humans find themselves outside, i.e., cast out to the environment, outside of these decisions. In such examples, human *agency*—acting on behalf of another, or providing a particular service—is being replaced by technologized agency. In categorizing the World Wide Web for instance, Google (to bring up one example) uses proprietary algorithmic robots (known as bots) to create a searchable database that then ranks users' search results based on their search queries. The structuring of the bot-generated entries contains the logic of how something will be "made searchable," though due to the complexity of the task, the interaction between bots and websites that are indexed must be unsupervised. But while the millions of *preindexed* search results give the illusion of choice, almost 90% of humans don't get past the top ten (Jansen & Spink, 2003). The whole process feels like a "search on the Internet" but it is actually a restricted human search of a technological presearch of the Internet: the "search of a search." Hence, this is a case of a human reacting to technological stimuli (i.e., an individual person reacting to the algorithmically generated search results intended to steer the person's behavior) rather than a technology reacting to human stimuli (i.e, a neutral search algorithm providing objective results to best serve a human using the technology).

Another well-known example comes from Amazon. The vast majority of prices are defined by *algorithms* in so far as Amazon vendors "use algorithmic pricing to ensure that they can automatically change their product prices based on a competitor" (Solon, 2011), with the result that vendors are being forced to engage in this practice for fear of losing out to the competition. Meanwhile, the algorithmic interactions between vendors carry the possibility of developing unpredictable consequences. Such algorithmic pricing on Amazon can be found in the example of the book entitled *The Making of a Fly* by evolutionary biologist Peter Lawrence. This book came to be priced at $23,698,655.93 (plus $3.99 shipping) as two sellers were using algorithms to adjust the price of the book in response to one another. It took 10 days for humans to notice and intervene to bring back the prices to normal levels; ironically, "normal levels" merely indicated a temporary human decision that would allow the continuation of algorithmic pricing.

Similar examples of human decisions and human reactions finding themselves outside the boundary of technologized decision-making, can also be seen in other domains. In law for instance, legal analysts are beginning to be replaced in complex cases by software that analyzes thousands of legal documents—proprietary e-discovery algorithms of software companies *pre*structure the defense/prosecution of a case by effectively preselecting an extremely limited subset of documents that will then be looked at and presented in court by humans (Markoff, 2011). The way that technology interferes with decisions is perhaps even starker in cases where personal liberties are at stake. For instance, in the United States, a large number of states use a variety of automated risk-assessment tools to decide (algorithmically) whether a prisoner should get parole (Kehl, Guo, & Kessler, 2017). The way this happens is both subtle and invisible (to an observer, i.e., a human judge who has no access to how the algorithm made the first decision attributing a risk score to a prisoner). This service is offered by a private, outsourced software company; essentially, algorithms scan prisoners' biographies to generate patterns that predict future criminal behavior (Walker, 2013). Thus, the algorithmic dependency upon which parole boards and judges rely is yet another obvious extraneous factor in judicial decisions (Danzigera & Avnaim-Pessoa, 2011). To make matters worse, an algorithmic bias against black defendants has been discovered when analyzing the same types of offenses; judges often rely on these risk assessments *without access to all the data that goes into the algorithmic calculations*. Risk scores for white defendants, in this case, were skewed toward lower-risk categories, but those for black defendants were almost static across different risk-score bands. For example, for the exact same background of arrests on two drug possession charges, Dylan Fugett (white, male) was given a low algorithmic-risk score of (3 out of 10), while Bernard Parker (African American, male) was given the highest possible algorithmic-risk score of 10 out of 10. After the release of both (over time), Fugett has been arrested three times on drug charges, while Parker has not reoffended (Angwin, Mattu, & Kirchner, 2016). In this case example, technology largely preconstructed the legal decision by designating a risk of reoffending and by risk-scoring prisoners. Human decision-making was thus shaped, influenced, preconstructed, and to some degree governed by an algorithmic backdrop against which human decisions were referred. Of course, somebody made design choices along the way, albeit bounded by the technology itself, regulation, and other factors. However, the visibility of the consequences of these choices will vary from one information system to another. In information systems that exhibit high degrees of complexity and interactions (such as the main example that we're exploring in this paper with algorithmic trading), the concept of the "IT artifact" mutates into that of a *system* that envelops its own consequences (on the basis of which new design choices are made).

By and large, it would be fair to say that in such cases, humans become the tool through which computerized decisions are voiced. In the case of autonomous driving, Google's self-driving car assumes full control of the vehicle, and while, in principle, it might appear as if human decisions can always override the technologized driving process, there are no strict criteria which humans could apply in the context of such a process (because the decisions of the vehicle necessarily depend on the unpredictable circumstances that emerge in its environment); in fact, the self-driving car caused its first crash when the human user was not certain whether he should have intervened (Davies, 2016). In the news domain, Bloomberg now uses automated sentiment analysis to provide financial news about companies in an automated way; this can further assist its customers in making financial decisions, "without having to consume the content" (Bloomberg, 2015). In the job market, many of the world's biggest companies use automated vetting software to screen résumés for particular job applications without even glancing at the résumés on the first round (Millar, 2012). Even in the judiciary, the discussion about whether computers would be better and fairer judges than humans has been around since 1977, and the algorithmic imposition of verdicts and fines has also been considered (D'Amato, 1977).

The examples listed above serve to illustrate (briefly) a two-step transition: the reduction of human agency, and the conditioning of humans based on technology acquiring systemic characteristics. Such characteristics are emergent and not reducible to any causal backdrop. Humans may conceive of their intervention as either not desirable (e.g., not financially beneficial or not influential in a system that develops itself algorithmically), or uncertain (e.g., humans are not sure if they should intervene to disrupt technologized decision-making, as such disruption could entail its own severe consequences). Humans might also avoid taking alternative actions that could be more beneficial to them, as a beneficial alternative might be masked due to an excessive reliance of human decisions on institutionalized technological stimuli. No matter the example, we can observe *the object (i.e. the unity) that is the distinction between humans/technology* differently. By simulating human decision-making in machines (e.g., autonomous driving), humans create a similar human/technology interaction within the domain of technology but are then subjected to its outcome, for which there needs to be another interaction that remains contingent on the former. Technology is expected to behave, search, invest, analyze, and decide as humans would, but the invisibility of how this is conducted in more complex systems, creates a series of contingencies. Of course, to establish the deeper second-order dynamics of these examples (in the tradition of second-order cybernetics), it would be necessary to delve into each example as a distinct case and explore the general recursive nature of human/technology interactions.

4 Theoretical Conditions for Considering Technology Itself as a System

In exploring a set of requirements/principles that systems theorizing must satisfy, Demetis and Lee (2016) consider general principles that are based on the founders of systems theory like Ludwig von Bertalanffy (1969), Kenneth Boulding (1956), James Grier Miller (1978) and Anatol Rapoport (1950). In their work, Demetis & Lee extract three principles from this body of systems theory and add another three from Niklas Luhmann (1995); they maintain that *systems theory* is particularly relevant for IS research and a rich theory to draw from (regardless of any epistemological differences amongst different scholars). In this context, we argue that systems theory lends itself to exposing the subtle emergence of technologized interferences in the relationship between humans and technology and will enable researchers (due to its abstract/generalized lexicon) to render novel cases through its concepts across several domains of application. In Table 1 below, we list the six requirements as they appear in Demetis & Lee (2016). We give each one a code (R_i) for when we use the requirements in this paper though it is worth stressing that we don't reflect on all six equally. Also, for the purposes of our discussion, we do not examine them in sequential order.

.

Table 1: Systems Requirements by Demetis and Lee (2016)

Requirement specified in Demetis and Lee (2016)	Code
Holism (the whole is more than the sum of the parts)	R_1
Goal seeking	R_2
Transformation process (of inputs into outputs)	R_3
Self-reference and autopoiesis	R_4
System/environment distinction	R_5
Communication	R_6

Our presentation of their requirements below is not a substitute for the original theoretical analysis—rather, it will be a discussion focused on targeting those theoretical conditions for thinking about a *system* comprised of technology. In other words, through different systems-theoretical principles, we can focus on the following question: *What systems theoretical principles would need to be adhered to, in order to consider technology itself as a system?*

One of the first requirements that Demetis & Lee (2016) stipulate in R_1 is the requirement to recognize that if we consider any *system* as a whole then "the whole is more than the sum of the parts." Regardless of whether we take technological artifacts in one sense (e.g., a combination of hardware/software) or another (e.g., sociotechnical assemblages), this requirement still applies. Whenever any elements become interrelated and interdependent, they form a more complex system; however: "One cannot deduce from complexity alone which relations among elements are realized" (Luhmann, 1995). Applied to technology at large, this systemic requirement demands that whatever technological "whole" is identified by an observer needs to be considered as a system that exhibits emergent properties. These properties remain irreducible to the whole's constituent elements.

Another important requirement (R_5) is based upon the fundamental distinction between *system* and *environment.* As stressed by Demetis & Lee (2016), this requirement is absolutely critical, as no system can be perceived without an environment (Luhmann, 1995). The environment should not be perceived as some type of residual category. Instead, "relationship to the environment is constitutive in system formation" (Luhmann, 1995, p. 176). By considering this fundamental principle that constitutes a pillar of systems theory and transcends all systems theorists (Hammond, 2003), we can pose an important question: if we consider technology as a *system* then what is its environment? One response would be that humans (also labeled by Niklas Luhmann as psychic systems) are in the environment outside of and around a *system of technology*; another response would be *other* technologies that do not interconnect directly with a specified (by an observer) *system* of technology. Another option of course would be a combination of the above—an *environment* wherein both humans/technology reside. While all possibilities can be considered, when we describe technology as a *system* in itself, we make the case that—via the systemic evolution of technology and its deep penetration in society—technology as *system* assimilates more functions within itself; in doing so, humans are cast out to the environment around the system of technology and cease to perform the function of decision-making (or are very limited in doing so). On occasion, the function of humans in the environment outside a *system of technology* becomes increasingly restricted to merely providing inputs through which technologized decision-making can continue uninterrupted. In this regard, the *system of technology* is the nondeterministic system that emerges from the complex interactions of networked IT artifacts, and which remains loosely coupled with an environment of humans, with the latter largely occupying the role of human artifacts in this system. It is a *system*, in the systems theoretical sense, so it is not simply IT or an IT artifact per se. Thus, where we consider technology as *system*, we have high degrees of technologized decision-making, weaker human/technology interactions, and stronger technology/technology interactions. As a whole, the behavior of the system of technology is irreducible to its constituent parts and thus linearity and causality do not apply.

At this point, it would be helpful to the reader if we acknowledged a difference—an extension from how Luhmann perceived technology. As we speak of a *system of technology*, we need to make it clear that Luhmann's depiction of technology was different and anchored on a set of a few principles that saw technology in a different light (it's also useful here to note that the types of phenomena that we refer to in this paper did not exist and Luhmann had not experienced the Internet as we know it today). For Luhmann, technology could be thought of as a "functioning simplification in the medium of causality" (Luhmann, 1993, p. 87). In plainer terms, Luhmann is simply recognizing that we have a reduction of an initial complexity, a simplification, which is channeled into technology. By considering causality as a medium, Luhmann depicts technology as a conduit of causal relations; as an enabler that converts (and reduces) an initial complexity into causality. In the case of technology, Luhmann considers the medium of *causality* as a representation of strict "operational continuity" (Moeller, 2006, p. 104). While we agree with Luhmann's depiction of technology, we find it more applicable for technology-oriented systems that are generally self-contained. But in a context of networked interactions (like in our example of algorithmic trading), we argue that causality is ultimately lost: causality dissipates at the level of the system (of technology) and controllability cannot be ensured. From this observing perspective, technological causality at the microlevel morphs into uncertainty at the macrolevel where it gives rise to emergent phenomena that cannot be attributed back to their original, and causal, representations (i.e., the whole of the behavior of IT artifacts is more than the sum of its parts). This leads to significant side effects. Reichel (2011, p. 106) frames this by saying that it is a convenient illusion, a beautiful lie to ensure the engineer, and in fact all human users of technology and their social communication about technology, that technology is manageable and the one last domain claiming predictability in our ever more contingent and insecure world" (Reichel, 2011, p. 106).

Opening the door to such a depiction of technology, Luhmann himself admitted that "the specious security of technology, based on repeatability and the control of defects, is a delusive one. This has consequences for the concept of risk" (Luhmann, 1990, p. 225) and this is exactly why he tackled technology in tandem with risk in his book chapter on "The Special Case of High Technology" (Luhmann, 1993, p. 83). Thus, an initial complexity is cast into causality at the microlevel through technology, and in turn, causal interactions give rise to risk. In a sense, the paradox of technology is the emergence of uncertainty and risk from within a series of streamlined (and forced) causalities (Angell & Demetis, 2010). Ultimately, the output of that paradox under conditions of reduced human agency leads us to a radical rethink of the role of humans in the interplay between humans and technology: the indeterminist-oriented, technological construction of humans by technology. While humans both act on, and are acted upon by technology, the combinatory effects of the deep penetration of technology in society and the emergence of a *systemic* character for technology, leads to circumstances where the latter condition prevails. When exactly this happens in a domain (like finance) might be difficult to ascertain; however, it is first important to recognize that it is indeed occurring and to explore the theoretical consequences of such a condition.

For humans, this is of course hard to accept and explore—in part, because it implies that humans simply *react* to what is now technologized decision-making and that, by and large, human "decisions" are secondary to systems-made decisions. At the very least, humans and what they believe to be their own decision-making, need to adapt to the emergence of systemic technological phenomena. Is human behavior patterning itself after, or simply being patterned by what the system of technology generates?

One implication leading from the primacy of the system/environment distinction is that the distinction can be replicated within the system (i.e., within any observed system, one can delineate further sub-*system/environment* distinctions). This is another way of saying that the system replicates that primary distinction (between system/environment) within itself. This idea of any *form* (defined by Luhmann as the *unity of a distinction)* affecting itself recursively is distilled in the concept of *re-entry*. Influenced by Bateson (1972), Luhmann makes the following remarks: "Accordingly, the re-entry of the form into the form—or of the distinction into the distinction, or of the difference between system and environment into the system—should be understood as referring to the same thing twice. The distinction re-enters the distinguished. This constitutes re-entry." (Luhmann, 2006, p. 54)

This type of re-entry is tightly connected with the concept of *self-reference* (R_4)*:* satisfying the systemic requirement for self-reference means identifying processes through which a system collects information about itself (and its own functioning), where this in turn can contribute to a change in its functioning. Through self-referential processes, certain systems (those that Luhmann calls autopoietic, a term he borrows from Maturana [1980]) continuously reproduce and maintain themselves. In this regard, considering technology as a *self-referential system* implies recognition of these dynamics that exhibit this form of re-entry; *technology* referring to *technology* is not a new phenomenon, but again, in the context of technologized decision-making, this acquires further significance. To the degree that technologized decisions become deeply embedded across different social systems, they elevate the complexity of the system of technology. This condition makes it harder for humans to gain visibility of the consequences of such systemic technologizing and, due to the excessive reliance of humans on technology, this reinforces the role reversal between the two. In a sense, the demand for further technologized decision-making ends up with humans augmenting the systemic character of technology further.

Another important systems requirement that follows Luhmann's systems theory involves the requirement for *communication* (R_6). For Luhmann, who reserves the term communication to indicate a broader mechanism (instead of only an act between human individuals), communication can be considered in the following triad form: announcement/utterance (*Mitteilung*), information (*Information*), and understanding (*Verstehen*). Also, in the event that human beings participate in the communication process, it would follow that the cognitive understanding (*Verstehen*) that would be developed by the recipient might not correspond to the intentionality of the individual conducting the utterance.

The reciprocity in communication, however, demands that whatever entity is receiving information following another entity's utterance will react based on its own (the former's) understanding. While we are typically used to thinking about these processes as human-based exchanges, they can be easily applied at a more abstract level. In fact, one of the pioneering theoretical implications of Luhmann's systems theory was the separation of the *communication* process from psychic systems (i.e., humans) and its consideration as a function of society (Luhmann, 1995, 2012). Keeping in mind that Luhmann himself had not addressed the issue of how the triad mentioned above could look in the context of technology and recognizing that there can never be a one-to-one correspondence between such contexts, we offer some additional thoughts.

In a system comprised of technology (rather than of humans), the hardest element of the triad to consider is *understanding (Verstehen*); however, that does not necessarily imply a cognitive understanding. In the case of a system of technology, substituting for human (or *cognitive*) understanding is what we label *technologized understanding*. Again, it is important to clarify that the concept of technologized understanding is not part of Luhmann's description but we find it important to introduce it as we seek to describe how the system of technology interferes with the recursive construction of human (cognitive) understanding. In our description, technologized understanding signals the a priori acceptance or rejection of any information, which would then prompt a computational response (another utterance) on the basis of a precoded algorithmic rationality. This demarcates the transition from (human) reflective understanding to preorganized understanding (that effectively collapses—or is reduced to—computerized decision-making). To describe this, we are considering the interplay between system/environment (i.e., technology /humans) and consider human artifacts to be characterized by a unique form of a *Verstehen of Verstehen* (echoing Foerster's 2002 book title "Understanding Understanding"). This demarcates a unique recursive conditioning of human, cognitive understanding that can be described as follows: (a) humans imbue with their *Verstehen* of a problem domain the modelling and the design of the conditions for algorithmic interactions that influence *computerized understanding* (this does not mean that computers "understand"; only that they render understanding in a computational frame). However,

(b) the complexity of networked interactions amongst scores of IT artifacts within a domain, dissipates the causality intended; complexity of interactions leads to nondeterministic outputs and these lead to risk. This is because elements can interconnect only by facing a necessary reduction in their own intrinsic complexity (and thus in their own potentiality). Without this necessary reduction in the intrinsic complexity of elements, connections cannot be established between them (Luhmann, 1995). In a sense, the dissipation of causality is a precondition for *computerized decision-making* to even exist and in turn, (c) the nondeterministic outputs that propagate from within the system of technology are disjoined from the intentions of human designers. Eventually, such outputs are (d) fed back into the environment where they reconstruct and reshape human *verstehen*. The very act of introducing technology, introduces a mechanism for computerized understanding, the output of which reconstitutes human *Verstehen*. The greater the reduction in human agency and the greater the complexity of interactions, the more masked the outputs and the stimuli that are spontaneously generated by technology. As we will see in our case, even in examples of—almost—catastrophic failure, it was not possible to attribute the problems to specific algorithms or traders, precisely because the *whole* (the system of technology) had become more than the sum of its limited parts (IT artifacts).

Also, human decisions are not just transferred across to computer decisions via the design of specific artifacts. The input from human designers is transformed (R_3) into system-technologized decisions/outputs through a complex nexus of technological interactions. While the goal-seeking (R_2) of specific algorithms may be perceived as fixed at the microlevel, at the macrolevel of the *system*, it becomes dynamic and carries the seeds of uncertainty. The very existence of unpredictable phenomena that emerge from such technological interactions is a testament to the fact that the correspondence between inputs and outputs in this context is nonlinear. As technology becomes responsible for “major systemic changes within the global financial sector . . . and as algorithms become ever more autonomous . . . we need a kind of ethical framework for developing algorithms” (Van Lier, 2016). Alas, even with a strict micromanagement of algorithmic developments, the very act of developing interconnections results in emergent unpredictability.

5 Case of The Dow Jones Index

First, let us clarify that it is beyond the scope of this paper to organize and/or deconstruct all the technologically oriented incidents behind the events of May 6, 2010 (what has become known as the “Flash Crash of 2:45”). Indeed, there remain several conflicting aspects on these events (Lhabitant & Gregoriou, 2015; Moosa, 2015; Sornette & Becke, 2011; Van Lier, 2016; Zaydlin, 2010) so we will focus here only on a handful of critical aspects that illustrate how the demand for increasing computerization is framing the interactions between humans and machines. In this regard, we would also like to highlight that the complexity of this case has also been attributed to the multifaceted role of algorithmic traders who make it almost impossible to deconstruct the case; this is acknowledged in several reports (see below).

The purpose of this section is to reflect on the case of the Flash Crash and connect some of its key aspects to a systems-based reconceptualization of technology based on the systemic principles put forward by Demetis & Lee (2016). This theoretical framing will allow us to consider the case of the Dow Jones index collapse as an example of technological domination—an example where the system of technology locates human artifacts in its environment, outside of and around the technology. The key sources that we will use in order to review the main findings of the Flash Crash include:

1. The report of the staffs of the U.S. Commodity Futures Trading Commission (CFTC) and Securities & Exchange Commission (SEC) to the Joint Advisory Committee on Emerging Regulatory Issues.
2. The UK report on “Crashes and High Frequency Trading” from the Government Office for Science (The Future of Computer Trading in Financial Markets—Foresight Driver Review—DR 7), and
3. An analysis for certified public accountants when advising investors entitled “Understanding the ‘Flash Crash’” that summarizes the basic characteristics of the Flash Crash in a succinct way (Betancourt, VanDenburgh, & Harmelink, 2011).

First, in order to set the scene, it is important that we reflect on the turmoil behind the specific events in their broader context. Uncertainty in May 2010 was already widespread in the market due to the possibility of a Greek government default on sovereign debt (Marouli, Caloghirou, & Giannini, 2015). Based on R5, we can think of this broader financial turmoil as the environmental stimulus to which the system of algorithmic traders would react. This negative market sentiment “was already affecting an increase in the price volatility of some individual securities” (CFTC/SEC, 2010, p. 1). This set in motion the following key events before the 2:45 p.m. crash (in chronological order):

1. A number of volatility pauses were triggered on the New York Stock Exchange (NYSE) around 1 p.m.; individual equities began to increase above average levels.
2. The S&P 500 volatility index rose by 22.5% by 2:30 p.m. Due to conditions of such volatility, investors moved their capital away from investments that were considered to be high-risk and towards safer options while the Dow Jones Industrial Average was down by 2.5% (due to selling pressure).
3. Buy-side liquidity in the "E-Mini" and the "SPY" (the E-mini S&P 500 futures contracts and the S&P 500 SPDR exchange traded fund) suffered 55% and 20% declines respectively.
4. Then, more critically, at about 2:32 p.m., a large fundamental trader (a mutual fund complex) "initiated a sell program to sell a total of 75,000 E-mini contracts (valued at approximately $4.1 billion) as a hedge to an existing equity position" (CFTC/SEC, 2010, p. 2).

The last step in the event is viewed as one of the critical triggers that led to the Flash Crash. But before we develop that further, we need to make an important observation in this context. Whereas traders can choose how much human judgment is involved when executing a trade (e.g., the trader can choose to enter orders manually in different time intervals, or indeed, outsource the process to a third party that will manage this process by conducting block-trades), over the years, the interference of human judgment came to be perceived (by humans themselves) as an obstacle. For example, Lewis (2014) remarks that during the crash of 1987 when the U.S. stock market fell sharply by 22.61%,

> *some Wall Street brokers, to avoid the orders their customers wanted to place to sell stocks, simply declined to pick up their phones . . . this time the authorities responded by changing the rules—making it easier for computers to do the jobs done by those imperfect people. The 1987 stock market crash set in motion a process—weak at first, stronger over the years—that has ended with computers entirely replacing the people*. (Lewis, 2014, p. 3)

The need for computerized decision-making was paired with the inexorable need for speed; the intensification of these dynamics led to Wall Street's "speed war." An example of that was the development of a superfast fiber-cable route between Chicago and New York by Daniel Spivey, just to shave "3 milliseconds off of the previous route of lowest latency" (Steiner, 2010). Speed matters; the fastest algorithm can exploit a large volume of minor discrepancies between markets and this is tantamount to "picking gold coins from the floor" (Steiner, 2010).

In such a context, where 825 miles of fiber are laid down through mountains, tunnels, and rivers, in the straightest line possible in order to shave off 3 milliseconds of trading, one can begin to contemplate how human beings are perceived (by the designers of algorithmic trading systems), considering the average reaction time (for the click of a button) is 215 milliseconds. Human beings are not just slow—they are (almost) redundant in such a market (though—as we shall see—they still serve the purpose of recalibrating (parts of) the technological system that is actually making the decisions). More recently, the search for more speed that would allow even faster versions of algorithmic trading to take place has led to large investments in microwave communications—with the goal of shaving off one extra millisecond in transacting (Westbrook, 2014).

It was in this context that a large fundamental trader initiated a sell program for the $4.1 billion trade; this was executed automatically, and thus *an algorithm was tasked to sell $4.1 billion*. In this case, the trader "chose to execute this sell program via a . . .'*Sell Algorithm'* that was programmed to feed orders . . . but without regard to price or time" (CFTC/SEC, 2010, pp. 2–3). The algorithm only took volume into consideration. But while the initial problems were indeed created by the algorithm of the fundamental trader, they were then "amplified by the strategic behavior of HFT [high-frequency trading]" (Sornette & Becke, 2011, p. 11). HFT is yet another name for algorithmic trading or black-box trading. The consequences of the amplification of the *Sell Algorithm* by yet other algorithms created a dynamic exchange between technologies of the same type. Based on R_3 this can be framed in the following general form: *Technology A* provides a stimulus for exciting *technology B*, and the output of B's operation(s), as feedback and input to *A*, recursively shape the environment of technological subsystem *A*.

One can conceptualize this situation through systems theorizing by means of considering alternative *observing perspectives* with which to explore the system/environment distinction (R_5). First, if the *Sell Algorithm* executed by the large fundamental trader is considered from the perspective of technological sub*system* A, then in its environment one can observe: (a) the totality of all other automated execution algorithms (let's call those *Algo(1), Algo(2), . . . Algo(n)*), and (b) the transaction outcomes from human traders (though as we saw in the introduction, the volume of their transactions is becoming more and more limited). Similarly, if we take the different *observing perspective* of another HFT, say *Algo(2),* which we define as our *system*, we would observe the *Sell Algorithm* in its environment, along with all other automated execution algorithms—Algo(1), Algo(3), . . . Algo(n)—as well as the human traders. Naturally, for fair competition reasons, the trading logic of any given *Algo(n)* is hidden from all other algorithms. Considering both of the mentioned *observing perspectives* in tandem at a metalevel (one that would be applicable for any automated execution), we have a *system* that is stimulated by the reactions it itself triggers upon its own environment.

Indeed, while the sell pressure established by the *Sell Algorithm* was initially absorbed by high-frequency traders (HFTs), at about 2:44 p.m., HFTs started to sell contracts aggressively. Then, the *Sell Algorithm* used by the large trader responded to the increased volume "by increasing the rate at which it was feeding the orders into the market, even though orders that it already sent to the market were arguably not yet fully absorbed" (CFTC/SEC, 2010, p. 3). This type of self-reference (R_4) can be portrayed as a "negative spiraling effect . . . (where) . . . HFT may have a destabilizing effect through its endogenous self-excitation nature within the (small) pool of participants" (Sornette & Becke, 2011, p. 11). This unsupervised self-excitation of technology at the level of the financial market is based upon a coupling between technology *and* technology, a condition that relegates humans to the external environment of that system. This has both significant and specific implications.

In the events that unfolded, the Dow Jones Industrial Average plunged 998.5 points. This became known as the Flash Crash of 2:45 (and while the index recovered some moments later, it wiped out value from several companies). After five months of investigations, the U.S. Securities and Exchange Commission (SEC), along with the Commodity Futures Trading Commission (CFTC), attributed the decline to the automated execution of orders despite the fact that no specific reason/trigger/algorithm could be identified for the event. In fact,

> *the exact reason or reasons for the so-called Flash Crash remain obscured by the mechanics of the electronic trading systems that execute millions of buy-and-sell orders during the course of a single trading day. Some initially blamed the crash on a "fat finger," (meaning a big individual investor) while others contended that essentially unregulated electronic trading platforms were the culprit. Others even questioned whether terrorists or hackers were behind the dramatic drop* (Betancourt et al., 2011, pp. 40–41).

The difficulty in identifying a specific cause renders the very idea of cause and effect (itself a nonsystemic idea) problematic in this context; this constitutes an irony, considering that the repeatable operations of technology usually allow the identification/back-tracking of effects. Of course, without recursive feedback, technology as a system would not exhibit such ramifications in the first place. In this regard, what would otherwise be conceptualized as cause and effect is better conceptualized as a web of back-and-forth impacts *distributed* amongst the complex interactions within the broader system of technology. There is no "error" in the individual technologies themselves (i.e., the algorithms); no bug that needs to be rectified. Thus, nonattribution of error to a single algorithm leads us to consider this as a *systemic phenomenon* that emerges out of the complex interaction of multiple automated execution technologies (R_1).

Of course, the Flash Crash would have been impossible at such a level without the complex interactions between algorithms. As the testimony of SEC Chairman Mary Schapiro confirmed to the U.S. Congress: "automated trading systems will follow their coded logic regardless of outcome, while human involvement likely would have prevented these orders from executing at absurd prices" (Schapiro, 2010, p. 7). The unpredictability with which automated algorithms feed off each other creates emergent conditions that can destabilize any system that technology itself penetrates. For the financial system that serves such an important function within the broader system of society, the implications are clear: Market outcomes are guided by computerized decisions that are executed algorithmically. While individual algorithms may reflect the general intentions of their designers, the algorithms as a whole find themselves within a far more complex environment (that they themselves help to constitute and recreate). In such conditions, algorithms feed off each other and—within the demands for millisecond transacting and communication (R_6)—create another version of the financial system where transacting decisions are executed in a technologized fashion; these lead to unpredictable consequences. Ultimately, this rearranges the roles of technology and humans.

For example, the chaotic behavior of the Flash Crash, triggered another algorithm, known as the *Stop Logic Functionality* of the Chicago Mercantile Exchange (CME). This was automatically executed in order to prevent any further price movements over an excessive range. By effectively pausing the market momentarily, this gives time to humans to recalibrate any parameters and relaunch the algorithms in the pit of algorithmic trading. The need to include *stop-and-pause* algorithms in order to contain any emergent uncertainties from the automated behavior of other buy/sell algorithms illustrates two things: first, how the "controllability" of technology is dependent on more technology, and, second, how restricted the role of human beings has become in the actual decision-making processes in the financial system. To a large degree, "*People no longer are responsible for what happens in the market, because computers make all the decisions*" (Lewis, 2014, p. 270). Whatever intentionality is imbued within specific algorithms by its designers, this dissipates quickly within a sea of complexity. Humans are merely there to "recalibrate" the relaunch of algorithms so that the latter can be pitted against each other in millions of millisecond transactions, the net result of which is both unpredictable and an emergent stimulus for further transacting.

Of course, once the demand for algorithmic decisions is firmly set, further complications can be considered. For example, new "crash algorithms will likely be developed to trade during periods of market stresses in order to profit from these periods" (Sornette & Becke, 2011, p. 4). Algorithms will also be developed that will attempt to exploit to the maximum the number of times they can offer/request something from the market. The patterns of some of these algorithms have been discovered on a few occasions, and what is rather astonishing is the speed with which orders can be placed. Below is a visualization of an algorithm (labeled the "Knife" by the company Nanex) where the algorithm transacts around 7 times every 60 milliseconds; in the image below, the dotted vertical lines indicate the 60ms intervals while the whole duration represented in the image below is 1.6 seconds (red: ask price, dark red: best ask, light green: bid, dark green: best bid). The analysis of the company on high-frequency trading found "cases where one exchange was sending an extremely high number of quotes for one stock in a single second: as high as 5,000 quotes in 1 second!" That's five quotes (for the same stock) every millisecond. As the company states: "Even more disturbing, there doesn't seem to be any economic justification for this" (Nanex, 2010). This level of logical detachment of computerized decision-making from the corresponding domain-specific decisions that humans would make becomes amplified in a system of technology; through the *system of* technology, the relationship between humans and technology is shaped asymmetrically, with the intensity of the influence of technology on humans shaped mostly by technologized decision-making.

Figure 1. Visualization of Algorithmic Formations During Trading (with kind permission from nanex.net)

Also, one could consider here algorithms that are designed to find the digital footprints of each other. By considering such algorithms that collaborate in specific market conditions, we could see the emergence of unfair *algorithmic* competition; and indeed, with the increasing attention that has been given to information security and cybersecurity breaches (Dhillon & Backhouse, 2000), automated financial trading may become subject to hackers or terrorists that will seek to create financial instability deliberately (e.g., by deactivating/modifying a stop algorithm). Whatever the destabilization effect, the speed of such transacting and the practical alienation of the human factor—relegated to the environment of the system of technology—can lead to an *algorithmic war* (similar to Amazon's algorithmic price escalations). To summarize, this involves algorithms that structure the market (by automated execution), antisystemic algorithms that could destabilize the system (e.g., by creating complexity in automated transacting and generating a financial crisis or other destabilization phenomena), crash algorithms that may exploit financial instability for profit, as well as the algorithmic response of the regulators to the crisis (through stop/pause algorithms). Due to the intrinsic complexity and the multiple entities involved in these cases, no single entity (human or algorithm) can monitor all financial interactions taking place at any given point in time. However, the ensemble of influences examined above highlight the necessary conditions for a system that maintains and sustains itself via the function of automation: the *system of technology*.

With our case in mind, humans are now cast out to the environment of such a system with restricted functions and reduced agency. This does not mean that humans are made redundant from our theorizing: only that the relationship between humans and technology needs to be reoriented and the intensity of this relationship considered in terms of both its consequences and its directionality. It is in this context, and inspired by Simon's (1996) indicia for artifacts, that we can now posit the following rudimentary principles for the human artifact. In each principle, we offer the systems theoretical context in which their interpretation can be supported. Consistent with the Popperian tradition that accepts paradoxes as unavoidable in any form of knowledge construction, we claim neither universality nor absolute truth in these principles. Under the conditions of a system of technology, we do hope that they are useful in reconceptualizing the role of humans within the distinction between humans/technology. It is also important to clarify that it is still humans that fill the role of human artifact in that it is a role that humans assimilate themselves into. This becomes activated in a spectrum of technologized decision-making, whereby human agency is greatly reduced (and humans become sidelined in their decision-making by offering mostly reactionary feedback that supports the uninterrupted operations of technology). We prompt other scholars to explore alternatives, reflect on their systems theoretical context, explore different systems approaches that could highlight other important aspects (e.g., complexity theory, boundary critique, system dynamics) and find new application contexts within which they can be enriched.

Table 2: Principles for The Human Artifact and Systemic Considerations

Principles for the human artifact	Systems theoretical context and implication
P1: Humans become artifacts shaped by the unintended consequences of complex interactions of IT artifacts (*role reversal*).	Human artifacts are in the environment surrounding a *system* of technology that emerges from the multiplicity of the complex interactions of IT artifacts. Humans as human artifacts continue to be loosely coupled with and affect the system of technology that emerges from such interactions. However, the combination of emergent stimuli from the system of technology and the excessive dependency of humanity on technology carves new boundaries for both technological and human agency. With no control over such a boundary, this could be labeled as "emergent indeterminism." In that space, human/technology interactions become weaker and technology/technology interactions become stronger; outcomes develop through high degrees of computerized decision-making.
P2: Human artifacts may imitate appearances of (natural) humans in the context of human decision-making, while lacking, in one or many respects, the reality of the latter (*reduced human agency*).	Human artifacts experience a shift in human decision-making. Reduction in human agency is a consequence of what occurs when interactions between IT artifacts become more than the sum of their parts and form complex networks within domains of society (e.g., the economy). Emergent phenomena, uncertainty and risk, become the background against which human decision-making is shaped. In the role of human artifacts, humans maintain a computationally contingent variant of human agency, with technology exercising a higher degree of non-causal influences on the boundary between the system of technology and its environment. It is the *system* of technology that influences human agency as it is irreducible to the lower-level properties and design of IT artifacts.
P3: Human artifacts are characterized in terms of functions and goals that are both limited in human agency and largely confined to support the uninterrupted operations of a system of technology (*heightened dependency of humans on technology*).	In contexts where technology is beginning to take over decision-making (e.g., algorithmic trading), the *system* of technology dominates the interaction between system/environment; human artifacts face an increasingly limited variety of decision-making that is both contingent and recursively shaped by the emergent phenomena and stimuli generated by the system of technology. The functions, goals, and adaptation of human artifacts are influenced by these systemic dynamics.
P4: Human artifacts can be understood as technology shaped and their roles are discussed and designed as responses to technologized imperatives *(technology-shaped agency)*.	The system of technology shapes what human artifacts must perform and act on with a reduced human agency in order to secure the uninterrupted operations of technologized decision-making.

By introducing the above indicia, we hope that the difference between the human artifact and the IT artifact will create a meaningful distinction through which the IS community can engage with those new phenomena that are characterized by a reorientation of human/computer interactions, a reduction in human agency, and the emergence of a system of technology. As the role of the human artifact is a consequence of the role reversal that we describe in our paper, we find it important to clarify the spectrum in which we see its applicability. Indeed, we do not claim that it would be equally applicable to all technologies. In order to describe this, we introduce the illustration below (Figure 2) and describe it in brief. We emphasize that this should be viewed as a visual-conceptual aid and not as a mathematical representation of thresholds or placements of different IS in the area of human-technological agency.

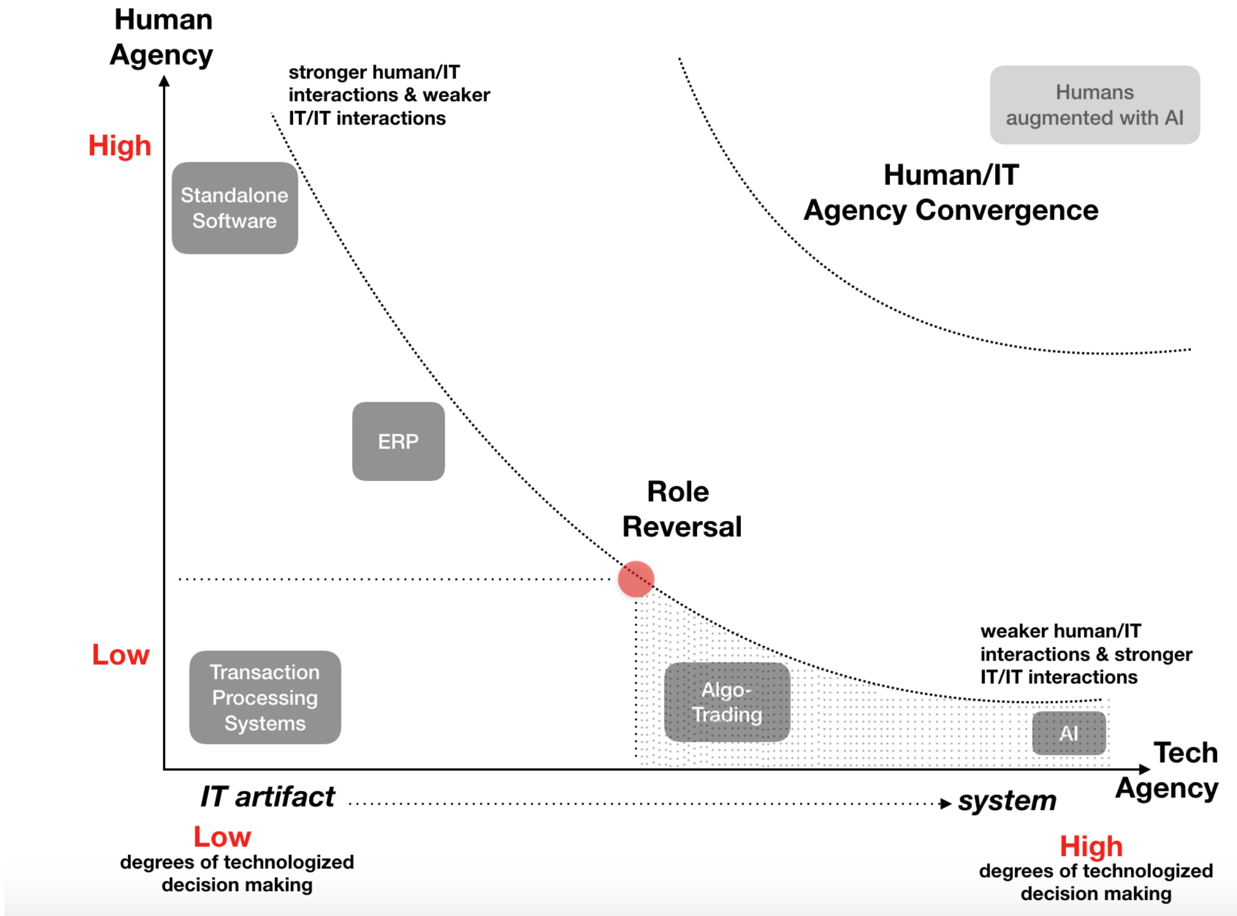

Figure 2. A Visual Metaphor for The Human Artifact In The Space Of Human/Technology Agency

At the x-axis where technological agency is expressed, we have the transition from low to high degrees of technologized decision-making while at the same time, we express this transition as part of the much broader trajectory that we describe: the emergence of a *system* of technology as we move to the right of the x-axis. In a similar manner, on the y-axis we have the spectrum of human agency. In cases where we have low degrees of technologized decision-making and high degrees of human agency, we find stronger human/technology interactions and weaker technology/technology interactions. The *combined effect* of a gradual reduction of human agency (lower part of y-axis) *and* the emergence of a system of technology (right part of x-axis) leads to a point (the role reversal) where technological agency becomes more significant and the role of the human artifact becomes dominant (with the shaded area in Figure 2 representing the area where the role of the human artifact applies within the human-technology agency space). In considering how different information systems would fit into that space, we have placed—indicatively—standalone software, ERP, algorithmic trading systems, and artificial intelligence, with the latter two placed in the area of where the human artifact would be prevalent. Also, in considering the two extreme scenarios of human/technology agency (low-low, high-high), we have transaction processing systems on one end (low human agency and low technological agency) while we have a more "futuristic" scenario on the other end of the spectrum (high human agency and high technological agency). Thus, in the space where both human and technological agency exhibit high degrees of decision-making, we find human/technology agencies converging—with one potential being humans augmented with AI and the development of a symbiotic agency (i.e., a hybrid agency). Of course, we recognize that organizations could place additional conditions onto human/technology agency expressions, thereby developing contingencies and moving IS into the human/technology agency space. Such conditions might be imposed by the structure of the organization, the mix of different technologies that are appropriated, the network of interactions, and so on.

The concept of the human artifact can also be anchored in existing scholarly work and explored in that context. As an example, we illustrate in brief, where the construct of the human artifact would fit with previous work by Lee et al. (2015) who describe the "IS artifact" as a *system* with three *subsystems*: (1) a technology artifact, (2) an information artifact, and (3) a social artifact. In their work, the authors argue that when these artifacts "are brought together and interact, they can come to form what we call an *IS artifact*" (p. 9). Thus, an IS artifact can be portrayed as the system that emerges from the interactions amongst its subsystems. In our paper, we are taking an observing perspective *within* the (metalevel) system of the IS artifact proposed by Lee et al. and thus, in our case, the human artifact would be a *sub-subsystem* of the social artifact, where the social artifact is a subsystem of the IS artifact system. Other sub-subsystems within the social artifact would include culture and social structure, with human beings occupying a double role in the social structure, one of human agency and another of human artifact.

The introduction of the role of human artifacts also has implications for humans as it signifies changes and unpredictability that cannot be harnessed by redesigning technological artifacts or reconfiguring specific algorithms (due to the irreducibility principle in the context of R_1). Humans are also greatly affected when they seek to interact within such domains. For example, an individual investor that is selecting a pension scheme may be unaware of the algorithmic trading strategies that would affect his own financial position, but ultimately she/he may be brought into the social artifact of the broader IS artifact that binds together technology, social, and information artifacts based on Lee et al.'s (2015) description. However, it is important to emphasize that this reconfiguration is not characterized by determinism. There is no regulation of this process as it is emergent. Outcomes are controlled neither by technology nor humans—they are the result of both strict and loose structural couplings across multiple dimensionalities of human/computer interactions that are in turn informed by other human/computer interactions: in other words, they are systemic. While we would be hesitant to label our position just yet, we could call it "emergent indeterminism" to differentiate it from both soft and hard determinism.

While we do describe a reduced role for human beings via the emergence of a system of technology (from the multiplicity of complex interactions of technology artifacts) and the emergence of the role of human artifacts, we actually stand against such an outcome. In calling attention to this role reversal, we are warning against it. Our concern is not only that "the specious security of technology, based on repeatability and the control of defects, is a delusive one" (Luhmann, 1990, p. 225), but that the role of human artifacts and the excessive reliance of society on technology, will create less controllable risks over time. The ensemble of these contingencies will circumvent human decision-making.

With these principles in mind, new technologies are designed as a *reaction* to how designers perceive the *system of technology* and wherever possible, designers code human decisions out of existence, thereby decreasing the actual involvement of human decision-making and accelerating the role reversal. These dynamics cannot be characterized by any form of determinism, as we've seen in our case. Whatever logic, controllability, and causality are injected into the technological domain, they dissipate quickly and are replaced by both uncertainty and unintended consequences. Technology has become systemic.

6 Discussion: From Artifact to System

As algorithmic domination in certain fields is changing the boundary between system/environment, between technology and humans, we feel there is a far greater role for the IS community to both pursue the consequences of these transitions and explore their implications within different organizational, societal, economic, and other contexts. Furthermore, with artificial intelligence on the rise, AI-based technologized decision-making is en route to reinforcing the systemic character of technology. Such transitions will require new theoretical constructs that should be able to open new horizons within the boundary between humans and technology; by introducing the construct of the human artifact, we hope to open up one such perspective in the context of a system of technology. By grounding this concept and the conditions that fuel it within systems theory, we hope we are contributing one part of the multiple contingent distinctions that will come into focus.

By conceptualizing technology as a *system*—with the help of systems theory—we must clarify that we view the optimization of any algorithms as a series of microefforts that will have marginal, *noncausal* effects in the behavior of the *emergent system of technology*. The phenomena that occur at the level of the system of technology cannot be reverse-engineered and attributed to specific starting points of algorithmic interference. Ultimately, through our case (the collapse of the Dow Jones index due to algorithmic trading) and our theoretical analysis, we offer a strong warning that there can be no controllability when an ensemble of IT artifacts acquires characteristics that are exhibited by emergent systems. Furthermore, this emergent system of technology has the capacity to highjack the function of important subsystems of society (like the economic system). Left uninterrupted, like in our example of algorithmic trading, technological operations may secure the continuous functioning of a subsystem of society (e.g., the economy) but this comes at a significant cost and with very high risks. This is also coupled by human decision-making becoming both marginalized and shaped by the nondeterministic effects of a system of technology. Thus, the emergence of unintended consequences and the unobservable complexities hidden behind algorithmic representations propel technology to a higher systemic level. Under specific circumstances (like those described in our case), human beings are restricted in the management of unintended consequences, an example of which is the intervention strategies chosen to relaunch the algorithms, the recalibration of trading thresholds based on human decision-making, the (human) design of new trading algorithms that will become part of the same nexus, general management decisions taken by humans and so on. While human beings, cast into the environment, can never be taken "out of the equation" as they are constitutive of the system of technology (as environments are always constitutive of the systems, based on R_5), they do become more loosely coupled with the IT artifacts, as routine decision-making shifts from humans to technology and human agency is reduced.

Despite our use of the Flash Crash incident, we are not making a value judgment about whether algorithmic trading is good or bad from an economic perspective. Indeed, there are scholars arguing that—overall—high-frequency trading by algorithms may even be beneficial for market liquidity (Anderson, Binner, Hagströmer, & Nilsson, 2015; Cliff & Northrop, 2012; Johnston, 2015). However, we would like to draw the attention of our readers to the fundamental imbalance in human/computer decision-making and highlight the interest that this poses for the field of IS. At a minimum, the conditions described above should allow us to reflect further on the transition from IT as an artifact (a tool shaped and used by humans to serve human ends) to IT as its own system (which, in turn, could regard humans as tools that maintain the systemic nature of technology). In that condition, technology gives rise to emergent phenomena and cannot be controlled in a causal way. Of course, this runs contrary to the design of technologies with a specified coded rationality.

The logic with which any given technological artifact interacts with other technological artifacts also requires reconsideration. In circumstances like those we have described in the previous section for the Flash Crash, one cannot deny that there is a high degree of complexity (not even a prolonged investigation could identify the "causes" as these are distributed and not linked to single entities). But there is another reading of complexity that could illuminate an additional aspect (we alluded to this in a previous section). Luhmann defines complexity

> *as a measure of the incapacity of a system to relate each element to every other one, be it in the system itself (system complexity) or in its environment (environmental complexity). . . .Complexity means the necessity of selective relations and, since relations specify what elements are possible within the system, complexity also means contingent elements. The analysis of complexity leads back to the notion of self-referential, self-organizing systems.* (Luhmann, 1983, p. 993).

Applied to the conditions that this view of complexity poses for technological interconnections, we can infer the following: *different technological artifacts must succumb to a restriction of their individual coded rationalities as a precondition to interconnect*. This serves to illustrate how the specificity of individual coded rationalities within any given IT artifact cannot be expressed due to the unavoidable restrictions imposed during interconnections. A very simple example of this is the following (this does not however capture the technologized complexity of more complex networks or entire function systems of society like the *economic system*): Suppose we have two different companies (A and B) that engage in algorithmic trading and both companies are the only two companies that want to sell/buy stock to/from another company X. We assume that A will not know what B's strategy would be and vice versa. By participating in algorithmic trading, A wants to sell the stock of company X and has set an algorithmic parameter between $20 and $10 (we're making the hypothesis here that price is the only criterion). Company B on the other hand has set an algorithmic parameter to buy the stock between $6 and $12. Both algorithms are executed so that their strategies are optimized (so the algorithm would first start exploring the sale of the stock at $20 before going down to $19.5, $19, and so on—one can inject here several other conditions like the time-frequency with which the drop would take place, the value of the drop in each step (say $0.10, $0.50, $1), etc). Similarly for B (starting from $6). But before even this process starts to take place for A and B, the subset of precoded rationalities, specifically the subset within the range between $12–$20 for A and $6–$10 for B, constitutes a nonexistent set for an observer who would have visibility of both. The establishment of allowed relations between A and B as a prerequisite for interconnecting, necessitates a restriction from the full spectrum of available possibilities for each one. Even in this simple example, seemingly well-defined thresholds that express precoded rationalities are facing restrictions based on their environment. The extent of unavoidable restrictions in element interconnection cannot be anticipated by the designers of the original systems of A and B. Despite the strict controls that may be imposed by designers, the very act of interconnection and communication, implies limitation.

Then, if we consider more realistic assumptions like: (a) dynamic ranges in price (say between \$x and \$y), which will not be fixed, or (b) algorithms that take input from their (uncertain and dynamic) environments in order to “determine” that (temporary \$x–\$y) range, we can see how millions of transactions and millisecond timeframes fuel systemic complexity. What is the role of designers in this case? What meaning does “controllability” acquire in the context of an “artifact?” Hence, the design of any technological artifact (such as an algorithm) with a specific coded rationality is simply the starting point through which that artifact will be allowed to partake in the complex nexus of algorithmic exchanges. Through those, all technologized trading algorithms “design” the market collectively and create an asymmetry between humans/technology; in those domains where technology has become more dominant in overtaking human decision-making, this implies a severe restriction of human agency, intentionality, participation, and decision-making. Consequently, the emergence of a system of technology is coupled with a decline in human agency and a rise of artificial humans (human artifacts).

This shift that we describe not only implies that “technologies create the ways in which people perceive reality” (Postman, 1993, p. 21); in taking decision-making away from humans, technologized decision-making within the context of a *system of technology* creates a reality that actually casts humans out to its environment. As such, human decision-making is becoming more and more restricted to a support/“tool-like” role that allows for the continuation of complex and invisible (at the level of the system) technologized decision-making. In mutating from an *artifact* (at the microlevel) to a *system* (at the macrolevel), technology carves new boundaries in the distinction between humans and technology. This presents new challenges but also opens up an important and novel domain for IS research.

7 A Research Agenda for The Future

Given the challenges that are presented in the context we describe, a number of additional questions can be raised. For the IS field that will find itself in a far more complex (and rapidly developing) space, where the reconfiguration between technology and humans will be anchored in a realignment between computerized decision-making and human decision-making, a series of questions will demand further exploration.

First, how can posthumanist approaches in IS research be adopted and encouraged? Given that there is a spectrum of interactions between humans and technology, how will management processes and the management of information systems be affected by the emergent system of technology? In exploring the new challenges that are posed by technologized decision-making in the context of the role reversal that we describe, we prompt IS scholars to look actively for cases that exhibit high degrees of technologized decision-making and to study the consequences these have for humans/human artifacts, organizations, and society at large within specific domains. These can come from different domains (e.g., e-commerce, law, finance, politics, science) and indeed, explorations from different domains would yield complementary insights about how technologized decision-making occupies the space of human/technology agency. Also, even though we do see emerging technologies (e.g., algorithmic trading, machine learning, artificial intelligence) as having a clearer influence in the role reversal that we describe, we also see the need for exploring the dynamics of the underlying digital infrastructures and networks in terms of how they support, develop, or restrict technologized decision-making. While we see the potential for many different approaches and applied research designs, given the nature of the phenomena at play, in-depth case studies would help create interesting insights into the conditions that prevail in the spectrum where the role of the human artifact applies.

Consistent with the concept of the IS artifact by Lee et al. (2015) and referring back to Figure 2, a series of further questions are important for further exploration:

a. How would different IT artifacts fit into the human/technology agency spectrum?

b. How would different IS artifacts recreate the human/technology agency spectrum and how would different organizational contexts affect the dynamics of the relationships described?

c. How can we enhance our understanding about what constitutes a stronger/weaker human/technology interaction and why?

d. What forms/types/classifications of technological agency can we delineate? What is their organizational and social impact/significance?

e. What other systems theoretical approaches could yield insights into these phenomena? While we use some key systems theoretical principles in this paper to discuss the emergence of a system of technology, we do think other systems theoretical analyses would lead to interesting paths of exploration (e.g., system dynamics, complexity theory). The very transition from artifact to system requires a rethinking of the systemic role of technology in society, and we see second-order cybernetics as opening many interesting paths of exploration in this regard, with an obvious one being the emergence of systemic risks from technology and pursuing a systemic description of the dynamics discussed above through elements and their relations (Luhmann, 1995).

f. We warn, however, that fixing ontological clarity for the concept of the human artifact would be as difficult as it is for the IT artifact. Exhaustive scholarly debate on the singularity of the concept of the human artifact would be counterproductive because we offer this concept in a particular light and context, prompting others to focus on what the *distinction* between the IT artifact/human artifact implies, since *identity* is only possible in comparison to *difference* (Bateson, 1972). By focusing on the *differences* of concepts as a *unit of analysis* instead of the concepts in isolation, a systems-theoretical discussion would further enable explorations into the contextual significance of the difference between the IT artifact and human artifact.

Of course, broader IS issues are important in the space that we describe. The ethics of IS are also important as phenomena, which, like the current pervasiveness of misinformation and fake news, will gradually become far more alarming and sophisticated, with the generation and communication of realistic content becoming entirely fictionate (e.g., AI-tools generating fake content and social bots spreading it). Our world is increasingly governed by technological systems that have deep effects in society; thus, understanding how power is delegated to such technological systems is critical (Jasanoff, 2016). It is essential that we understand the relationships between IS design in complex systems and early warning signs (where possible) that would assist the interventions of human agency when necessary. But where lies responsibility for complex IS failures when the systemic character of technology does not allow us to pinpoint exactly what/who is responsible? Can we develop approaches where human responsibility is connected to algorithmic responsibility, and how can the latter be dissected and even connected to designers? In other words, how can the new *system* of technology be governed and what is the role of IS design in that context?

As technologized decision-making continues to change the landscape of human/technological agency, the systemic character of technology elevates the need to address the fundamental shifts that we describe in this paper. The consequences of the role reversal between humans and IT create new challenges but also open new and exciting possibilities for research.

References

Alter, S. (2015). The concept of "IT artifact" has outlived its usefulness and should be retired now. *Information Systems Journal*, *25*(1), 47–60.

Anderson, R. G., Binner, J. M., Hagströmer, B., & Nilsson, B. (2015). Liquidity: Systematic liquidity, commonality, and high-frequency trading. In G. Gregoriou (Ed.), *Handbook of high frequency trading* (pp. 197–214). Amsterdam: Elsevier.

Angell, I. O., & Demetis, D. S. (2010). *Science's first mistake : Delusions in pursuit of theory*. London: Bloomsbury Academic.

Angwin, J., Mattu, S., & Kirchner, L. (2016, May). Machine bias: There's software used across the country to predict future criminals. And it's biased against blacks. *Propublica*. Retrieved from https://www.propublica.org/article/machine-bias- risk-assessments-in-criminal-sentencing

Bateson, G. (1972). *Steps to an ecology of mind*. New York, NY: Ballantine.

Benbasat, I., & Zmud, R. (2003). The identity crisis within the IS discipline: Defining and communicating the discipline's core properties. *MIS Quarterly*, *27*(2), 183–194.

Bertalanffy, L. (1969). *General system theory*. New York, NY: Braziller.

Betancourt, L., VanDenburgh, W., & Harmelink, P. (2011). Understanding the "Flash Crash." *The CPA Journal*, *81*(1), 40–45

Bloomberg. (2015). Sentiment analysis of financial news and social media. Retrieved from http://www.bloomberglabs.com/data-science/projects/sentiment-analysis-financial-news-social-media/

Boulding, K. E. (1956). General systems theory: The skeleton of science. *Management Science*, *2*(3), 197–208.

CFTC/SEC. (2010). Findings regarding the market events of May 6, 2010: Report of the staffs of the CFTC and SEC to the joint advisory committee on emerging regulatory issues. Retrieved from https://www.sec.gov/news/studies/2010/marketevents-report.pdf

Chaboud, A., Chiquoine, B., Hjalmarsson, E., & Vega, C. (2009). Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market. *International Finance Discussion Papers*, 980, 1–44.

Cliff, D., & Northrop, L. (2012). The global financial markets: An ultra-large-scale systems perspective. In R. Calinescu & D. Garlan (Eds.), *Large-scale complex IT systems. Development, operation and management*. Springer: Berlin.

D'Amato, A. (1977). Can/should computers replace judges? *Georgia Law Review, 11* (1277 Can/should computers replace judges?1281).

Danzigera, S., & Avnaim-Pessoa, L. (2011). Extraneous factors in judicial decisions. *Proceedings of the National Academy of Sciences of the United States of America*, 108 (17).

Davies, A. (2016, February). Google's Self-Driving Car Caused Its First Crash. *Wired*. Retrieved from https://www.wired.com/2016/02/googles-self-driving-car-may-caused-first-crash/

Demetis, D. S., & Lee, A. S. (2016). Crafting theory to satisfy the requirements of systems science. *Information and Organization*, 26(4), 116-126.

Dhillon, G., & Backhouse, J. (2000). Information System Security Management in the New Millennium. *Communications of the ACM*, 43(7), 125–128.

Foerster, H. von. (1971). Responsibilities of Competence. *Journal of Cybernetics*, 2(2), 1–6.

Foerster, H. von. (2002). *Understanding Understanding: Essays on cybernetics and cognition* (pp. 29–70). Berlin: Springer.

Hammond, D. (2003). *The science of synthesis: Exploring the social implications of general systems theory*. Boulder: University Press of Colorado.

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, *28*(1), 75–105.

Jansen, B., & Spink, A. (2003). *An analysis of web documents retrieved and viewed*. Paper presented at the 4th International Conference on Internet Computing. Las Vegas, Nevada, U.S.A.

Jasanoff, S. (2016). *The ethics of invention: Technology and the human future*. New York, NY: Norton.

Johnston, J. (2015). Information events, big data, and the flash crash. In U. Ekman, J. D, Bolter, L. Diaz, M. Sondergaard, & M. Engberg (Eds.), *Ubiquitous computing, complexity and culture*, (pp. 378–392). Abingdon, U.K.: Routledge.

Kehl, D., Guo, P., & Kessler, S. (2017). Algorithms in the criminal justice system: Assessing the use of risk assessments in sentencing. Retrieved from http://nrs.harvard.edu/urn-3:HUL.InstRepos:33746041

Lee, A., Thomas, M., & Baskerville, R. (2015). Going back to basics in design science: From the information technology artifact to the information systems artifact. *Information Systems Journal*, *25*(1), 5–21.

Lewis, M. (2014). *Flash boys: Cracking the money code*. New York, NY: Penguin.

Lhabitant, F. S., & Gregoriou, G. N. (2015). High-frequency trading: Past, present, and future. In G. Gregoriou (Ed.), *Handbook of high frequency trading* (pp. 155–166). Amsterdam: Elsevier.

Luhmann, N. (1983). Insistence on systems theory: perspectives from germany. *Social Forces*, *61*(4), 987–998.

Luhmann, N. (1990). Technology, environment and social risk: A systems perspective. *Industrial Crisis Quarterly*, *4*, 223–231.

Luhmann, N. (1993). *Risk: A sociological theory*. New Brunswick, NJ: Transaction.

Luhmann, N. (1995). *Social systems*. Stanford, CA: Stanford University Press.

Luhmann, N. (2002). *Theories of distinction: Redescribing the descriptions of modernity*. Stanford, CA: Stanford University Press.

Luhmann, N. (2006). System as difference. *Organization*, *13*(1), 37–57.

Luhmann, N. (2012). *Introduction to systems theory*. Cambridge, U.K.: Polity.

March, S., & Smith, G. (1995). Design and natural science research on information technology. *Decision Support Systems*, *15*, 251–266.

Markoff, J. (2011, March 4). Armies of expensive lawyers, replaced by cheaper software. *The New York Times.* Retrieved from http://www.nytimes.com/2011/03/05/science/05legal.html?_r=0

Marouli, A., Caloghirou, Y., & Giannini, E. (2015). Non-performing debt recovery: Effects of the Greek crisis. *International Journal of Banking, Accounting and Finance*, *6(*1), 21–36.

Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and cognition : The realization of the living*. Boston, MA: Reidel.

Millar, M. (2012, November 9). Beating the recruitment machines. *BBC News.* Retrieved from https://www.bbc.com/news/business-20255387

Miller, J. G. (1978). *Living systems*. New York, NY: McGraw-Hill.

Moeller, H.-G. (2006). *Luhmann explained*. Peru, IL: Carus.

Moosa, I. (2015). The regulation of high-frequency trading: A pragmatic view. *Journal of Banking Regulation*, *16*(1), 72–88.

Nanex. (2010). Analysis of the “Flash Crash”— Part 4: Quote stuffing. Retrieved from http://www.nanex.net/FlashCrashFinal/FlashCrashSummary.html

Orlikowski, W., & Iacono, S. (2001). Research commentary: Desperately seeking the “IT” in IT research: A call to theorizing the IT artifact. *Information Systems Research*, *12*(2), 121–134.

Popper, K. (2002). *The Logic of Scientific Discovery*. Abingdon, U.K.: Routledge.

Postman, N. (1993). *Technopoly : The surrender of culture to technology*. New York, NY: Vintage.

Rapoport, A. (1950). *Science and the Goals of Man: A Study in Semantic Orientation*. New York, NY: Harper.

Reichel, A. (2011). Technology as system: towards an autopoietic theory of technology. *International Journal of Innovation and Sustainable Development*, *5*(2/3), 105–118.

Schapiro, M. (2010). Testimony concerning the severe market disruption on May 6, 2010. Retrieved from http://www.sec.gov/news/testimony/2010/ts051110mls.htm

Silver, M., & Markus, L. (2013). Conceptualizing the sociotechnical artifact. *Systems, Signs & Actions*, *7*(1), 82–89.

Simon, H. (1996). *The sciences of the artificial*. Cambridge, MA: Massachusetts Institute of Technology Press.

Solon, O. (2011, April). How a book about flies came to be priced $24 million on Amazon. *Wired.* Retrieved from http://www.wired.com/2011/04/amazon-flies-24-million/

Sornette, D., & Becke, S. von der. (2011). *Crashes and high frequency trading: An evaluation of risks posed by high-speed algorithmic trading*. UK Government Foresight Project: The Future of Computer Trading in Financial Markets. Retrieved from https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/289016/11-1226-dr7-crashes-and-high-frequency-trading.pdf

Steiner, C. (2010, September). Wall Street’s speed war. *Forbes.* Retrieved from http://www.forbes.com/forbes/2010/0927/outfront-netscape-jim-barksdale-daniel-spivey-wall-street-speed-war.html

Van Lier, B. (2016). From high frequency trading to self-organizing moral machines. International *Journal of Technoethics, 7*(1), 34–50.

Walker, J. (2013, October 11). State parole boards use software to decide which inmates to release programs look at prisoners' biographies for patterns that predict future crime. *The Wall Street Journal*. Retrieved from https://www.wsj.com /articles/state-parole-boards-use-software-to-decide-which-inmates-to-release-1381542427

Westbrook, J. (2014). High-frequency traders find microwaves suit their need for speed. Retrieved from http://www.bloomberg.com/news/articles/ 2014-07-24/high-frequency-traders-find-microwaves-suit-their-need-for-speed

Zaydlin, A. (2010, June). Regulating in response to the "flash crash." *Economist*, 395(8686). Available at https://www.economist.com/letters/2010/06/10/ on-malaysia-wynne-godley-the-flash-crash-toll-roads-denmark-artificial-life-gays-in-africa-german

About the Authors

Dionysios S. Demetis is a lecturer in management systems at the Hull University Business School. He is also vice-chair of the Humber Business Resilience Forum, which promotes a secure Humber region by enhancing the cybersecurity of its businesses. Dr. Demetis holds a PhD on anti–money laundering and information systems from the London School of Economics and Political Science. He is the author of two books, one on the philosophy of science and the intrinsic paradoxes of knowledge creation (entitled *Science's First Mistake: Delusions in Pursuit of Theory*, published by Bloomsbury), and another on the systems theoretical deconstruction of Anti-Money Laundering and the role of information systems in it (entitled *Technology and Anti-Money Laundering: a systems theory and risk-based approach*, published by Edward Elgar). He has authored numerous other publications, and his research on anti–money laundering has been featured in the United Nations bibliography and has been funded by the European Commission.

Allen S. Lee is professor emeritus of information systems at Virginia Commonwealth University. He served as associate dean at both Virginia Commonwealth University and McGill University, as editor in chief of *MIS Quarterly*, and as a founding senior editor of *MIS Quarterly Executive*. His research program has involved identifying basic lessons from the philosophy and history of science and extending them in the information systems discipline to show not only how qualitative research can be done rigorously, but also how quantitative research equally needs to live up to the requirements of science. He is a fellow of the Association for Information Systems, a member of the Circle of Compadres of the Information Systems Doctoral Students Association of the KPMG PhD Project, and a founder of Chinese American Professors of Information Systems. In 2015, he received the LEO Award for "lifetime exceptional achievement in information systems" from the Association for Information Systems.