

Chapter 8: Tacit Knowledge and the Future of Work Debate by Michael Piore¹

Not for Circulation Without Permission of the Author

This note is an attempt to understand “tacit knowledge” and how a recognition of it changes the way we think about work and labor. It is part of a larger project which focuses on the policy paradigms that direct contemporary economic policy, shaping the way we evaluate the performance of the economy and what we believe we can do to improve it (Piore 2018). The central concern of that project is the paradigm which has come to dominate current debates, that I have come to think of as the Silicon Valley consensus (Alloa and Soufron 2019; Moore and Tall 2016). Its essence is captured by the mantra “innovation and entrepreneurship in the knowledge economy.” In the debate about the future of work, this has led to a kind of technological determinism, a belief that the job structure will be bifurcated between jobs requiring high levels of formal education and those jobs that are unskilled, typically service work (Autor 2014). Mid-levels jobs will disappear and the jobs at the bottom of the skill hierarchy will be forced increasingly to compete with labor in low-wage developing countries (Autor, Dorn and Hanson 2013; Autor and Dorn 2013). This view seems to explain, if not justify, the increasingly skewed distribution of income that has become a central concern of policymakers and politicians in the United States. It also seems to explain the decline of manufacturing and the hollowing out of the industrial heartland where the electorate abandoned its traditional mooring in the Democratic Party to elect Trump in the 2016 Presidential election.

Each of the components of the Silicon Valley mantra is, however, open to question. The thrust of the broader project is to try to unpack them, and to understand them separately and in relation to each other (Autor, Dorn and Hanson 2016). Here we are directly concerned with the concept of knowledge but, as will become apparent, tacit knowledge also changes the way in which we conceive of innovation and of the role of entrepreneurship in bringing about innovation.

The advent of AI has shifted the debate about the future of work since this project was initially conceived. But that debate is still characterized by a technological determinism and a commitment to the notion that the economy is driven by innovations conceived in Silicon Valley by especially creative and entrepreneurial individuals.

What do we mean by tacit knowledge

¹ David W. Skinner Professor of Political Economy, Emeritus, Department of Economics, Massachusetts Institute of Technology. Text originally prepared for presentation at the 2019 Society for the Advancement of Socio-Economics SASE Annual Meetings, New York, June 27, 2019.

The dictionary definition of tacit knowledge is: Knowledge that you do not get from being taught, or from books, etc., but from personal experience.² The term is generally attributed to Michael Polanyi in his 1966 essay (Polanyi 1966) who characterized it by the idea that we know more than we can tell, things we know but cannot explain; in a sense we cannot even say how we know them. It is knowledge or understanding that we acquire by being in the world. Polanyi's examples range from knowing how to ride a bicycle to our ability to recognize a face. But it is also organizational or professional knowledge. Significantly, when it is recognized and cultivated in prestigious occupations, for example in medicine or law, we use the term "clinical knowledge". What is unclear, at least to me, is whether tacit knowledge and formal knowledge are in any sense the same thing. The notion of a body of shared knowledge that cannot be communicated through verbal communication emerges in work divided among members of a larger team, in, for example, software programming or fashion design (Piore 2004). Tacit knowledge is uncoded; but one of the puzzles which surrounds it is whether it *could* be codified. And even if it could be codified to the extent that it could be formally communicated and transmitted, can it be scientifically understood?

The term was recently introduced into the debate about the future of work and the ability of technology to replace human labor by David Autor in his paper on "Polanyi's Paradox" (Autor 2014). He argues that the capacity to generate and maintain tacit knowledge is a uniquely human trait for which machines would never be a complete substitute. But the argument here is that the implications of tacit knowledge are potentially much more fundamental, and call into question the analytical apparatus we are accustomed to using for thinking about and formulating an answer, and that Autor's conclusions, while not necessarily wrong, are premature.

My own understanding of tacit knowledge is rooted in observations of the introduction of new production equipment on the manufacturing shop floor. I first began to look at this when working on my Ph.D. thesis (further back in time than I really care to admit) (Doeringer and Piore 1971), but surprisingly what I saw then is not so very different from what I have seen in trips to the field over the intervening years. We thought then that new technology was embodied in equipment and that the equipment came with a blueprint that the plant personnel, or the vendor, used to install and operate it (the original plan for my thesis was to collect and compare these blueprints and the manning tables that we thought would accompany them, in plants producing similar products over time). But in fact it turned out that the actual installation was more ad hoc than that. New equipment was typically installed on the plant floor by the engineers who designed it or the vendor who sold it. They worked with the plant's own skilled maintenance craftspeople and sometimes with the operators as well to get the equipment up and running, and in this sense the design engineers (or the vendor sales reps) constituted the first generation of operators. But in the process of installation and initial operation they gradually transferred the work, both operation and maintenance, to workers who had less formal education and training.

This was in effect a deskilling of the jobs, at least in the sense that skill is defined in terms of formal education and training (a point to which I will return to shortly). But what appears to be the case is that the workers who take over understand the work in a different way from the engineers who designed it or the vendors who sold it. Their understanding is tacit knowledge in the sense that the workers have a great deal of trouble articulating what they are doing; they

²See, <https://dictionary.cambridge.org/us/dictionary/english/tacit-knowledge>.

cannot express their understanding in words. Because it is different from the engineers' understanding, we could say that they are inventing a new way of doing the job, that the transfer from the engineers to the operators is a form of innovation, which involves the creation of new knowledge, new at least for the workers who come to possess it, but quite possibly totally new to the world. But it is often not simply an innovation in knowledge; the workers often change the equipment itself in order to make it easier to operate or change the way in which they relate to their colleagues in the production process facilitating the flow of product and material as production progresses in time and in physical space. Thus, the innovation in understanding extends to innovative changes in the physical and social environment in which the work takes place. In this sense, the understanding that develops in this largely ad hoc process is an example of what is called in the literature "embodied cognition" (Shapiro 2011): Our understanding is embedded in the physical and social environment in which we are operating (Menary 2010). This process of developing new ways of working and manipulating tools and equipment does not necessarily stop once the equipment is installed but continues with the equipment in use. Arrow calls it learning by doing, and asserts that it is responsible for what he, following Lundberg (1961), calls the "Horndal effect", named after a Swedish iron works where over a 15-year period productivity increased at an annual average of two percent with no additional capital investment (Arrow 1962).

Ethnographic studies involving the close observation of work as it is performed and of the transfer of skill in craft and professional training reveal similar characteristics: ad hoc learning, nonverbal communication, embodied cognition, and skill development and/or skill transfer as a continuous process.³ They reveal a process, in other words, in which learning tends to involve, and, for the outside observers at least, morph into invention and innovation. For example, Matthew Beane describes the evolution of work in a surgical operating room that he observed over a period of several years, following medical students as they joined the work team and then progressed through internships and residencies to become accredited surgeons. He describes a continuous redistribution of roles and responsibilities as new recruits first begin as essentially passive students and then move to more active roles, assisting and then taking over procedures themselves and ultimately becoming teachers and team leaders in their own right, and the way in which that process is disrupted as it tries to accommodate the shift to robotic surgical technology (Beane 2019).

Natasha Iskander and Nichola Lowe describe the way in which immigrant workers on construction sites in Philadelphia figure out how to perform the work assigned to them by inexperienced owners, adapting and improvising techniques they brought from Mexico to the tools and materials they find available in Philadelphia, trying to imitate experienced US workers observed surreptitiously and at a distance, but drawing on their own experience in Mexico to figure out what they are doing (Iskander and Lowe 2010, 2013). Iskander goes on to observe Asian immigrants working in Qatar on monumental buildings with unique and untested designs, functioning as a team at enormous heights, unable to speak a common language, but nonetheless managing to figure out not only how to put the building together but to do so working on barely protective scaffolding. Here they survive only by supporting each other physically and moving in

³ For example, Iskander's work (Iskander 2021) on construction in Philadelphia, Atlanta and Qatar; and Beane (2019) on the training of urology surgeons.

coordination so as to maintain the balance of the team and prevent any member from falling off (Iskander 2017, 2021).

As noted earlier, somehow this same phenomenon of interdependence managed through non-verbal communication appears to underlie the way members of a large software development project, each working on a separate piece of the project, all of which must ultimately be brought together to form the final program, come to share a common vision of the overall architecture (Brooks 1995; McBreen 2002; and Piore 2004). Still another example of the importance of nonverbal communication involves the ethos of the agents of a large bureaucratic agency who have considerable discretion but learn to exercise it in a common way, treating similar cases in the same way their colleagues do (Piore and Schrank 2018; Orr 1996; and Barley and Orr 1997).

All of these traits suggest that the way in which tacit knowledge merges and evolves is like a language—indeed one could almost say that language *is* tacit knowledge. A language can be codified by a grammar and a vocabulary, but it cannot be reduced to these formal characteristics, and native speakers do not typically learn a language by studying its formal structure. Language, moreover, is social, it is shared by a community of people. It is contextual. And, it mixes together verbal and nonverbal communication, for example “I don’t know what I mean until I see how you react to what I am saying” (Jackendoff 2007).

Task-based analysis and the limits of the division of labor in the analysis

These features of tacit knowledge are not readily captured through the conventional framework that we use for thinking about job structure. That framework thus limits our ability to recognize and understand tacit knowledge and the way in which it emerges and evolves over time. This is because the framework attempts to break up jobs (and the capacities required to perform them) into a series of elementary components and to focus on those components separately and in isolation from each other and from the social and physical environment in which they are found. In the case of jobs, those components are “tasks”; worker capacities to perform these tasks are generally termed “skills”. Structural problems in the labor market are then understood in terms of the match, or mismatch, between skills and tasks. The debate about the impact of technological change on jobs focuses on the ability of new technologies to perform individual tasks (Brynjolfsson, Mitchell and Rock 2018; McAfee and Cummings 2014). (Although policy proposals almost invariably have focused on the provision of the ability to perform these tasks through education and training, as opposed to efforts to change the tasks that the workforce is required to perform.)

This was a major advantage when technological progress was focused on mechanization, as it was for most of the last two centuries, a point to which we will return shortly. The problem is that it is a way of looking at production which ignores the interrelation among the tasks and the possibility that the efficiency of production may depend upon that interrelation and the way in which that evolves over time in any given production process through the transfer of skills among workers in a learning process, or as the process changes through an interrelated set of changes over time. This is the heart of what seems to occur in the development and application of tacit knowledge. The importance of focusing on the interrelationship among operations was the basic lesson of Japanese management, which American producers “discovered” in the latter

part of the last century (Nonaka and Takeuchi 1995). The failure to take that into account is, to take another example, what produces bugs in software programming. But the interrelationship among tasks, and indeed the domain of knowledge which they may require, also limits the capacity to understand technological change which has, particularly in recent years, involved integration across domains of knowledge. The cellular telephone, for example, involved the integration across radio and telephone technologies that initially were totally separate and independent. Much of medical innovation involves the integration of biology and mechanical engineering (Lester and Piore 2004).

The focus on skills, defined as the tasks the worker is capable of performing, has much the same limitation as the focus on tasks to define the job. It is static. It particularly fails to capture the way in which workers' skills tend to change over time, reflecting the essentially continuous process of learning on the job. Because it neglects tacit knowledge, moreover, it inevitably tends to overemphasize institutions which certify (and presumably provide) skill. And because tacit knowledge tends to be embedded in the physical and social environment, it turns attention away from the processes which change that environment, limiting it and rendering some capacities obsolete even as they remain critical to work, but also presumably enhancing the ability to generate and expand other worker capacities (tacitly).

The task-based approach to innovation and work goes back at least to Adam Smith and his famous pin factory. The process of economic growth and development for Smith entailed the division of labor in which the work of the master pin maker was transformed into the pin factory where "one worker pulled the wire, the second worker cut the wire, a third worker pointed the pin, a fourth headed it". The tradition of thinking about technological change in this way was reinforced in modern American manufacturing by Frederick Taylor, the discipline of industrial engineering which Taylor founded, and Ford's automobile assembly line. In analysis but also in practice, it leads to a focus on individual tasks in isolation from each other.

In Smith's own writing, the transition to the pin factory is linked to a broader argument about specialization as the engine of economic growth and development. But his own writings and those of a host of analysts that followed raise questions as to whether the extreme specialization associated with the pin factory is actually efficient (since the tasks are so trivial and repetitive that workers lose interest in the job and no longer pay attention to what they are doing).

An alternative explanation that circumvents this issue is suggested by Karl Marx in an argument which connects the task-based organization of work to mechanization. The underlying difference between craft production and modern mechanization, Marx argues, is that the former uses tools that are general purpose and in use are an extension of the human being and the logic of human action, whereas in mechanization this relationship is reversed and the human becomes an extension of the machine. The difference is such that, for example, one could not mechanize the pin makers' work directly because by its very nature it takes its meaning from the human beings who perform the work and consume its products. One could not even imagine a single machine that could do that. But once you break up the work into a series of separate tasks, and focus on each task in isolation, it is possible to conceive of a specific machine for each task—a machine which pulls the wire, a machine which cuts the wire, etc.—and to mechanize the operation in this way.

The distinction drawn by Marx is given wider resonance by Roman Jakobson, an extremely influential linguist but not, so far as I know, recognized in economics. Jakobson distinguishes two dimensions of human cognition: *metaphor* and *metonymy* (Jakobson 1971). Metaphor separates an object from its environment in order to identify similarities to other objects in different surroundings. Metonymy looks at the relationship among objects in their own environment. Each of these dimensions of cognition is located in a different part of the brain. The difference emerges when the brain is impaired by a stroke which affects a person's capacity for speech. Such strokes are, Jakobson argued, of two very different kinds. In one type of stroke, the victim loses the capacity to name an object with a single word, but can still say that same word in a whole phrase or sentence. In the second type of stroke, the victim can name objects in isolation but cannot construct sentences containing that word. Understood in these terms, the limit of task-based analysis is that it focuses on metaphor to the exclusion of metonymy. It takes the object out of one context in order to assimilate it into a different context, but is unable to deal with meaning that is dependent on context. Thus, when cognition is embedded in the social and physical environment, as appears to be the case in tacit knowledge, the task-based approach virtually excludes it from its purview.

Jerome Bruner draws a similar and evidently related distinction between two ways of organizing knowledge. One of these is the way we organize formal, or "scientific", knowledge in terms of chains of causal relations. The other is narrative in which we organize knowledge in terms of a story which arranges events and understands the relationship among them in a temporal sequence (Bruner 1990, 1991). There is a small literature on narrative understanding of work which draws on the kind of close, ethnographic observation of the work process upon which our understanding of tacit knowledge is based (Barley and Orr 1997). It reveals many of the same characteristics as tacit knowledge: The way in which it grows organically out of the work process, and the way in which it is socially embedded and evolves as it is passed on from one generation of workers to another, leading to innovations in the way in which work is performed.

But narratives are clearly not tacit. Indeed, Bruner stresses the way in which narratives both draw upon and presuppose a shared language and at the same time the language itself evolves and changes through the narratives. In the work context, one can think of this evolution as innovation in Arrow's learning by doing (Arrow 1962). It is unclear how one should treat the relationship between tacit knowledge and narrative. Is narrative a third form of knowledge? Or should one think of tacit and narrative knowledge as part of a broader category, alternative to formal knowledge?

Despite the apparent contradiction in terms, it seems particularly appropriate to draw upon linguistic theory for an understanding of what is at stake in tacit knowledge and how it might differ from formal knowledge because, as already suggested, language embodies and exemplifies the difference between them. A language has a formal structure reflected in a grammar and a vocabulary which can be abstracted and "scientifically" described, independently of the culture and context in which it is used. But a language cannot be reduced to these formal dimensions. And the formal rules of a language do not seem to have much to do with the way in which native speakers learn to speak it. This is in part because there are so many exceptions to the rules, but also in part because the language is continually evolving through use. It quickly outgrows the

formal structure that is captured at any moment of time. But, more important, to participate in the society to which a language belongs, one must align oneself with that evolutionary process and participate in it, and the formal rules do not tell us how to do that (Jackendoff 2007).

Conclusions

The implications of this view of tacit knowledge for analysis and policy are too far reaching to explore in a brief note of this kind. What we have suggested is less a fully developed theory of tacit knowledge than a critique of the concept of formal knowledge and of the associated concepts of skill, innovation and entrepreneurship around which the Silicon Valley consensus has been built (but also much of contemporary thinking about economic progress more broadly). The chief limit of conventional analysis is that it leads us to think in terms of stocks – accumulations of knowledge or of skill, for example. Innovation then becomes a way of adding to the stock of knowledge, and education and training a way of adding to the stock of skill. The notion of entrepreneurship suggests that expansion of these stocks is a deliberate process which is generated by particular and identifiable individuals (who, not incidentally, then appear to deserve the rewards associated with it). Tacit knowledge, by contrast, seem to emerge automatically through the work process, and innovation and skills emerge from it, part and parcel of and virtually inseparable from work. That process is moreover not one of personal or individual effort but is rather a social process, i.e., one in which people interact and cooperate with each other. What we think of as economic progress is often the unexpected byproduct of the work process. This is better characterized by Arrow's notion of technological change as "learning-by-doing," in contrast to the kind of radical change evoked by the image of the independent entrepreneur. The promise is that were we to better understand the nature of tacit knowledge, the process through which it is generated could be guided and directed to produce better economic and social outcomes—more equal economic opportunity, or a more equal distribution of income and social welfare.

This does not directly address the question at the heart of the future of work debate: Will/Can technology completely replace human agents and render them irrelevant to productive processes? This, however, may be the wrong question. If, as we have argued, the impact of technology on work is contextual, the result of a social process, it will depend very much on the characteristics of the environment in which it is introduced, the distribution of income in that environment, for example, or the nature of the job hierarchy, the boundaries of the organizational unit, the way in which jobs are defined and responsibility distributed. In sum, the different characteristics that, in the conventional analytical approach, we treat as separate and distinct from the technology and view basically as institutional features are those that we can try to adjust to mitigate the impact of technology. Whereas in the end these features of the work environment are inseparable from the technology and with it will determine how work evolves. Thus, if it were meaningful to separate out technology as an independent variable and assess its impact—and we stress again that this may not be possible—it would turn out that the impact will vary from one workplace to another depending on these other characteristics.

But the other problem in addressing the question is that the very nature of the technology whose impact we are attempting to predict is unclear, especially in the domain that distinguishes tacit from formal knowledge. Initially the dominant approach to AI was to replicate (or try to

replicate) *formal* decision-making, and hence to leave tacit knowledge, almost by definition, beyond its reach. But as the capacity of machines to accumulate, store and manipulate data has exploded in recent years, a very different approach has emerged. This second approach essentially uses the machine to look for correlations between the problem at hand and a vast library of similar problems that have already been solved. The focus is on statistical fit, circumventing the need to understand the underlying relationships that produce it (e.g., Zittrain 2019; Stewart 2019). This would seem to make tacit knowledge more accessible to machine learning than formal knowledge, rather than less so. What makes this conjecture particularly plausible is that, to the very limited extent that analysts have tried to model tacit knowledge and the way it is applied to practical problems, this is precisely the way the human brain is assumed to work (Abell 2011).

None of this is to suggest that formal knowledge and the mechanisms that create it and distribute it are irrelevant. Nor that what we have called tacit knowledge exhausts the different ways in which know-how is generated. We have suggested that narrative understanding is important in certain craft jobs, and substitutes for causal *analysis*, but narratives are certainly not tacit. We have not moreover explored the question of whether tacit knowledge can be formalized or vice versa. But if tacit knowledge is an important part of the process of economic growth and development, and if indeed it is embodied in social settings and emerges through a social process, what is missing most in the debates about work and innovation is an understanding of the forms which that process takes and the implications for economic and social welfare. The need for greater attention to these social forms is probably the most important lesson to emerge here. This has already begun to happen around the edges of organizational studies (Lorenz and Kraemer-Mbula 2019; Lam 2000; and Mintzberg 1980).

Bibliography

- Abell, Peter. 2011. "Singular Mechanisms and Bayesian Narratives." In Pierre Demeulenaere (Ed.) *Analytical Sociology and Social Mechanisms*. Cambridge, UK: Cambridge University Press, pp. 121–135.
- Alloa, Emmanuel and Jean-Baptiste Soufron (Issue Editors). 2019. "L'idéologie de la Silicon Valley." *Revue Esprit*, May.
- Arrow, Kenneth J. 1962. "The Economic Implications of Learning by Doing." *Review of Economic Studies* 29 (3): 155–73.
- Autor, David. 2014. "Polanyi's Paradox and the Shape of Employment Growth." NBER Working Paper No. 20485.
- Autor, David H., and David Dorn. 2013. "The Growth of Low Skill Service Jobs and the Polarization of the U.S. Labor Market." *American Economic Review* 103 (5), 1553–1597.
- Autor, David H., David Dorn and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103 (6), 2121–2168.
- Autor, David H., David Dorn and Gordon H. Hanson. 2016. "The China Shock: Learning from Labor Market Adjustment to Large Changes in Trade." *Annual Review of Economics* 8: 205–240.
- Barley, Stephen R. and Julian Orr (Eds.). 1997. *Between Craft and Science: Technical Work in U.S. Settings*. Ithaca, NY: Cornell University Press.
- Beane, Matthew. 2019. "Shadow Learning: Building Robotic Surgical Skill When Approved Means Fail." *Administrative Science Quarterly* 64 (1): 87–123.
- Brooks, Frederick P. 1995. *The Mythical Man-Month: Essays on Software Engineering*. Reading, MA: Addison-Wesley.
- Bruner, Jerome. 1990. *Acts of Meaning*. Cambridge, MA: Harvard University Press.
- Bruner, Jerome. 1991. "The Narrative Construction of Reality." *Critical Inquiry* 18 (1): 1–21.
- Brynjolfsson, Erik, Tom Mitchell and Daniel Rock. 2018. "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" *AEA Papers and Proceedings* 108: 43–47.
- Doeringer, Peter and Michael Piore. 1971. *Internal Labor Markets and Manpower Analysis*. Lexington, MA: Heath.
- Iskander, Natasha. 2017. "The Right to Have 'Society in the Bones': The Skill and Bodies of Male Workers in Qatar." *Women's Studies Quarterly* 45 (3/4): 234–244.
- Iskander, Natasha. 2021. *Does Skill Make Us Human? Migrant Workers in 21st-Century Qatar and Beyond*. Princeton, New Jersey: Princeton University Press.
- Iskander, Natasha, and Nichola Lowe. 2010. "Hidden Talent: Tacit Skill Formation and Labor Market Incorporation of Latino Immigrants in the United States." *Journal of Planning Education and Research* 30 (2): 132–46.
- Iskander, Natasha, and Nichola Lowe. 2013. "Building Job Quality from the Inside-Out: Mexican Immigrants, Skills, and Jobs in the Construction Industry." *ILR Review* 66 (4): 785–807.
- Jackendoff, Ray. 2007. *Language, Consciousness, Culture: Essays on Mental Structure*. Cambridge, MA: MIT Press.

- Jakobson, Roman. 1971. "Two Aspects of Language and Two Types of Aphasic Disturbances." in Jakobson, Roman and Moris Halle (Eds.) *Fundamentals of Language*. 2nd ed., rev. Berlin: Mouton de Gruyter, pp. 69-96.
- Lam, Alice. 2000. "Tacit Knowledge, Organizational Learning and Societal Institutions: An Integrated Framework." *Organization Studies* 21 (3): 487–513.
- Lester, Richard A. and Michael J. Piore. 2004. *Innovation – The Missing Dimension*. Cambridge, MA: Harvard University Press.
- Lorenz, Edward and Erika Kraemer-Mbula. 2019. "Organizational design, forms of knowledge and the automation debate." Paper prepared for the EMAEE Workshop, Sussex University, 3-5 June 2019.
- Lundberg, Erik. 1961. *Productivity and Profitability. Studies of the Role of Capital in the Swedish Economy*. Stockholm: Norstedt.
- McAfee, Andrew and Jeff Cummings. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: WW Norton & Company.
- McBreen, Peter. 2002. *Software Craftsmanship: The New Imperative*. Boston: Addison-Wesley.
- Menary, Richard (Ed.) 2010. *The Extended Mind*. Cambridge MA: MIT Press.
- Mintzberg, Henry. 1980. "Structure in 5's: A Synthesis of the Research on Organization Design." *Management Science* 26 (3): 322-41.
- Moore, Matt and Kelly Tall. 2016. "The Silicon Valley Consensus and What It Means for Information Professionals." *Online Currents* 30 (2): 89-95.
- Nonaka, Ikujiro and Hirotaka Takeuchi. 1995. *The Knowledge Creating Company: How Japanese Companies Create the Dynamics of Innovation*. New York, NY: Oxford University Press.
- Orr, Julian. 1996. *Talking About Machines: An Ethnography of a Modern Job*. Ithaca, NY: Cornell University Press.
- Piore, Michael J. 2004. "The Limits of the Division of Labor in Design and the Prospects for Off-Shore Software Development in Mexico." presented at 16th Annual Meeting on the Society for the Advancement of Socio-Economics, Washington DC.
- Piore, Michael J. 2018. "Economic Policy in the Time of Reactionary Populism." *Issues in Science and Technology*, 34 (3): 24-29.
- Piore, Michael J, and Andrew Schrank. 2018. *Root-Cause Regulation: Protecting Work and Workers in the Twenty-First Century*. Cambridge, MA: Harvard University Press.
- Polanyi, Michael. 1966. *The Tacit Dimension*. Garden City, NY: Doubleday.
- Shapiro, Lawrence. 2011. *Embodied Cognition*. New York: Routledge.
- Stewart, Matthew. 2019. "The Actual Difference between Statistics and Machine Learning." Medium.com, accessed April 23, 2025: <https://medium.com/data-science/the-actual-difference-between-statistics-and-machine-learning-64b49f07ea3>
- Zittrain Jonathan. 2019. "Intellectual Debt: With Great Power Comes Great Ignorance." Medium.com, accessed April 23, 2025: <https://medium.com/berkman-klein-center/from-technical-debt-to-intellectual-debt-in-ai-e05ac56a502c>