Article

S|S|S

Social Studies of Science
2015, Vol. 45(2) 242–269

DOI: 10.1177/0306312714564912
sss.sagepub.com
SAGE

Thinking science with thinking machines: The multiple realities of basic and applied knowledge in a research border zone

Steve G Hoffman
Department of Sociology, University at Buffalo, The State University of New York, NY, USA

Abstract
Some scholars dismiss the distinction between basic and applied science as passé, yet substantive assumptions about this boundary remain obdurate in research policy, popular rhetoric, the sociology and philosophy of science, and, indeed, at the level of bench practice. In this article, I draw on a multiple ontology framework to provide a more stable affirmation of a constructivist position in science and technology studies that cannot be reduced to a matter of competing perspectives on a single reality. The analysis is grounded in ethnographic research in the border zone of Artificial Intelligence science. I translate in-situ moments in which members of neighboring but differently situated labs engage in three distinct repertoires that render the reality of basic and applied science: *partitioning, flipping*, and *collapsing*. While the essences of scientific objects are nowhere to be found, the boundary between basic and applied is neither illusion nor mere propaganda. Instead, distinctions among scientific knowledge are made real as a matter of course.

Keywords
applied science, Artificial Intelligence, basic science, boundaries, ontology, scientific practice

The distinction between basic and applied science is often invoked casually, as if everyone is drawing upon the same criteria. Yet, voluminous scholarship contains little consensus as to what, if anything, this distinction maps. How can we, as analysts of scientific knowledge

Corresponding author:
Steve G Hoffman, Department of Sociology, University at Buffalo, The State University of New York, 430 Park Hall, Buffalo, NY 14260-4140, USA.
Email: sgh@buffalo.edu

production, reconcile the multiple ways of talking about the boundary between basic and applied science? Given the increasingly elaborate entanglements between industry and academia that have developed over the past few decades (Etzkowitz et al., 2000; Rhoten and Powell, 2010; Shore and McLauchlan, 2012; Vallas and Kleinman, 2008), it is worthwhile to raise again the old question of how this boundary is drawn.

It will prove helpful to simplify the distinction between basic and applied science to two main ways of apprehending what in that distinction is real: one view sees the distinction in the *substance* of knowledge itself; another treats it as an *artifact* of some exterior and/or a co-constitutive condition.[1] I argue that while it is a clear misrecognition to reduce the distinction to a set of invariant characteristics, it is also a category mistake to decide that scientific objects are artifice, illusion, language game, pure discourse, or mere propaganda (for a compatible critique, see Latour, 1993 [1991]: 5–8). Recent formulations of relational theory (Mol, 2002; Strathern, 2005; see also Emirbayer and Mische, 1998; Sewell, 1992; Vallas and Cummins, 2014) and multiple ontology (see the June 2013 issue of this journal), together with close attention to the demarcation of symbolic boundaries (Calvert, 2006; Gieryn, 1983, 1999; see also Abbott, 2001; Lamont, 2009), provide ample resources that can be brought to bear on this issue.

Research scientists produce basic and applied research in multiple and sometimes contradictory ways. What is considered basic in one scenario might be considered applied in another (which can give the distinction the impression of being mere illusion or rhetoric). The boundary can be collapsed altogether (suggesting it is passé or even never existed). I argue that no single form of the boundary is true or correct in any singular, trans-locational way. Similar to the multiple versions of atherosclerosis, ostensibly the same disease, in a single hospital (Mol, 2002), the boundary between basic and applied knowledge multiplies in everyday contexts. While this position affirms the idea that the boundary is constructed in the sense of having been enacted by a set of activities and rhetoric, multiple ontology shifts analytic attention away from competing perspectives and toward the phenomenological performance of contrasting realities (Schütz, 1945). I focus on how scientists make multiple versions of basic and applied science cohere rather than deducing what this boundary really is, or should be, in some ultimate way. A multiple ontology framework moves analysis beyond the representationalism that muddles the research policy debates over university–industry entanglements.

It is for this same reason, however, that I do not begin with a formal definition of the most important characteristics of basic and applied knowledge. Instead, I show how the boundary is performed. I find that three repertoires[2] render distinct versions of the basic–applied boundary: *partitioning, flipping*, and *collapsing*. Basic and applied sciences are made substantive, artifactual, and indeed, both or neither, in practical action. Academic Artificial Intelligence (AI) provides an opportune empirical case, since it is an area of research that has long straddled the overlapping borders of science and technical application (Crevier, 1993; Edwards, 1996: 294–302; for a proto-history, see Galison, 1996). AI originated as an academic discipline with contradictory visions that, since its founding moment, continue to influence how its practitioners invoke the difference between basic and applied science. On one hand, AI was inspired by universalistic inquiry into the nature of human-level intelligence. On the other hand, it was grounded in the practicalities of symbolic programming and computer engineering. This tension is not a necessary or natural aspect of computer science, per se, but it is a deep source of strain that is highly unlikely to be resolved within university labs devoted to AI. That tension certainly cannot be resolved in this article. Instead, I explain how members of two academic AI labs actively manage this tension between knowledge-for-scientific-understanding and technologies-for-use. In this way, I demonstrate that the basic and applied sciences do not reside in essences of scientific objects, but that certain aspects of technoscience are made more substantial than others.

Getting clear on or steering clear of the basic and applied boundary

In his 1945 report to President Truman on the state of post-World War II (WWII) science, Vannevar Bush, Director of the US Office of Scientific Research and Development, argued that 'pure' research is the fundamental driver of new technologies: 'New products and new processes … are founded on new principles and new conceptions … painstakingly developed by research in the purest realms of science' (p. 14). US university administrators, scientists, and policy allies reified this relationship between basic and applied science through a distribution of federal research dollars. A similar boundary is found in the rhetoric of English educational reformer John Tyndall:

> Let the self-styled practical man look to those from the fecundity of whose thought he, and thousands like him, have sprung into existence. Were they inspired in their first inquiries by the calculations of utility? Not one of the them. (quoted in Gieryn, 1999: 54)

Tyndall's advocacy, like Bush's, turned on the notion that the development of path-breaking concepts and principles is necessarily distinct from and drives research-for-use.

Over the past three decades, policymakers and university administrators in the United States and elsewhere have sought to link academic research and economic growth more directly (Gieger and Sá, 2009; Popp Berman, 2012; Rhoten and Powell, 2010). This reorientation has created formal and informal pressures on academics to produce knowledge products that have industrial relevance, are 'translational', are patent-worthy, span disciplines, or can be moved quickly from 'bench-to-bedside' (Hoffman, 2011; Moore et al., 2011; Slaughter and Rhoades, 2004). A good deal of management scholarship identifies conditions that ease the transfer of academic knowledge to industry (e.g. Rhoten and Powell, 2007; Thursby and Thursby, 2010). Critical scholars, in contrast, view this trend as a worrisome decline in the importance of basic science (e.g. Krimsky, 2003). Voluminous empirical research uses loose measures of scientific productivity to assess whether applied work has pushed out basic research (e.g. Van Looy et al., 2006). Across the practical, policy, and scholarly work in research policy, varied as it may be, is an underlying assumption that the difference between basic and applied science is 'within' knowledge itself.

This *substantive* conceptualization of the boundary between basic and applied can also be found in the philosophy and sociology of science, where we find periodic attempts to discern how substantive differences map onto how generative a piece of scientific research is (Stinchcombe, 2001), onto disciplinary-level variations (Frank and Gabler, 2006), onto modes of inquiry (Sintonen, 1990), or onto multi-dimensional layering of such variables (Lomnitz and Cházaro, 1999). Consider how Stinchcombe (2001) turns the dichotomy into a continuous measure of generativity:

> Fundamental notions in science play a central role because they generate solutions to a wide variety of problems … structural members with the right size, materials, and angle to the force of gravity are principal determinants of structural soundness. The size, materials, and angle of electrical conduit, on the other hand, make very little difference to anything else. (p. 160)

In this architectural metaphor, the weight-bearing bits of knowledge are basic because additional knowledge depends upon their foundation. Other analysts offer new classificatory schemes (Gibbons et al., 1994; Stokes, 1997) or pose the shift as a rise in a hybrid order (Lam, 2010; Murray, 2010; Owen-Smith, 2003). Whether conceived as binary, continuous, or hybrid, however, the boundary between basic and applied science is still treated as a substantive essence.

Science and Technology Studies (STS) scholars in the constructivist tradition, in contrast, have conceptualized the basic–applied boundary as a historical and rhetorical process (Latour, 1987; Shapin, 2008; for a genealogy, see Calvert, 2006: esp. 201–203). Gieryn (1983) pointed toward this *artifactual* position in his analysis of scientific boundaries:

> Demarcation is as much a practical problem for scientists as an analytical problem for sociologists and philosophers. Descriptions of science as distinctively truthful, useful, objective or rational may best be analyzed as ideologies: incomplete and ambiguous images of science nevertheless useful for scientists' pursuit of authority and material resources. (pp. 792–793)

While the focus on scientific boundaries has been highly productive, the emphasis on how the line separating basic and applied sciences has varied historically has created a tendency to dismiss the boundary as passé (e.g. Edgerton, 2004; Forman, 2007; less problematically, Sturdy, 2007). Edgerton (2004) provides the quintessential example of this, although arguments like Forman's (2007), that science (writ large) became subsumed under technology around 1980 rest on similarly shaky ground. Edgerton (2004) argues that the distinction was never more than academic propaganda aimed at keeping federal dollars flowing:

> 'The linear model' is usually taken to be something like the following: 'basic' or 'fundamental', 'pure' or 'undirected', scientific research in the main source of technical innovation the process of innovation is a sequential one, by which discoveries arising in such research are developed in a sequence through applied research, development and so on, to production … 'the linear model' not only did not exist, but it could not exist as an elaborated model … It is always, however, something that was surpassed, criticized, to be moved beyond … 'The linear model' is a term of art without a history. (pp. 32–34)

Here, there is no substance to separate basic and applied research at all. Not now or ever.

The position that the basic–applied boundary has only ever served as propaganda or 'art' may seem to resonate with a central tenet of constructivist STS, namely that science does not reveal timeless truths but that facticity is produced within assemblages and networks (Bijker, 1995; Collins, 1985; Latour, 1987; Pinch and Bijker, 1987). It echoes the metaphysics of pragmatism too, in which ontology is located not in the object but in its use (Dewey, 1922). However, by suggesting that the boundary 'did not exist', 'could not exist', and is a 'term of art without a history', the claim subtly glides from process to a static ontology of reflection.

Multiple ontology is useful for pushing STS past this trap. The ontological turn in STS is compatible with the development of relational theory more generally (e.g. Alexander, 1982; Emirbayer and Mische, 1998; Vallas and Cummins, 2014). This work focuses on how scientific objects, boundaries, and things more generally are coordinated in such a way that they are made politically, professionally, and very often personally important within particular webs of practice. A relational approach insists that what is real must be 'built up' into categorical distinctions (Blumer, 1969), yet their coordination suggests that the real might always be otherwise. As Latour (1987) argues, this is an accomplishment to follow, not decide.

The case of AI as a border zone

Academic AI straddles a research border zone.[3] Since its inception as a professional field in the mid-20th century, it has been a science of sentience caught between the crosscurrents of autotelic knowledge and knowledge for an externalized purpose (McCarthy et al., 1955). The abstract question 'What is intelligence?' must concede to the practicalities of computer code. As such, AI science rests on a double-layered ambiguity. First, there is no consensus over the nature of intelligence, cognition, decision-making, or analogical reasoning. Second, there is no consensus over how to reproduce this vague referent in machinery. Entrenched within this double ambiguity, AI scientists forge ahead with local bets. Open-ended questions about human-level reasoning are transformed into something that peers might accept as good enough, such as, 'How can we get this computer program to recognize a drawing of a wheelbarrow if we draw it in two slightly different ways?' AI science spans this unresolvable but fecund border zone. As Alfred North Whitehead (1938) might have formulated it, the assumptions embedded in an AI technology are 'always degenerating into philosophic generality'. Ambiguity lingers 'just on the edge of consciousness' (p. 70), bracketed by concessions to precision but never eliminated.

I observed two academic AI labs at a private research university in the American Midwest. Both are part of the same Department of Computer Science in a top-ranked research university. The university is not explicitly technology-driven like MIT or Carnegie Mellon, but recently placed near the top of the annual rankings of US public and private universities by royalty income as measured by the Association of University Technology Managers. The Department is divided into several semi-autonomous research groups, two of which identify as AI labs. Lab spaces are separated from department spaces, although they are located in the same building. This separation provides a measure of privacy, enabling members to draw cross-lab comparisons without offending their neighbors with an unflattering evaluation. Each lab has an operating budget subsidized by the department as well as by external grants. Student committees are assigned according to lab membership, with a lab head serving as chair.

Three types of data comprise the empirical reportage: observational, interview, and textual. All three are mobilized toward a holist description oriented to analytic induction (Geertz, 1973; Marcus, 1986). I conducted observations over 3 years, between 2004 and 2007, and I remain in occasional contact with both labs. I attended and audio-recorded over 200 lab-wide and/or project meetings, along with system demos and informal conversations. I was actively involved in project meetings and also read drafts of lab papers to offer grammatical or clarification advice. My relative technical incompetence enabled me to ask a lot of simple questions that an experienced AI practitioner might consider embarrassingly rudimentary. After about 1 year of fieldwork, I developed 'interactional expertise' (Collins and Evans, 2002) within the AI subfields of the two labs – a level of expertise sufficient to understand the basic assumptions, algorithms, and formalisms embedded in lab systems and to be conversant on the broader field. Field notes comprise a synthesis of my daily logs of lab activities and audio recordings of conversations and events. I find audio recording indispensible to my own recall, which has all the typical fallibilities of the human species (Lehrer, 2011) and ethnographers (Fine, 1993). Observations are rendered in the past tense; both labs, however, remain active.

I also conducted in-depth interviews with all members of the two labs. I used interview data to help triangulate ethnographic observations, asking individual members to reflect on their experiences, occurrences, and viewpoints. Interviews were open-ended, although I used a guide to cover issues such as professional history, description of research trajectory, broader professional networks, how interviewees evaluated good AI research, and the appropriate boundary between basic and applied research. Interviews lasted anywhere from 1 to 4 hours, often conducted over several days. I conducted follow-up interviews with all lab members in my third year of fieldwork. I also collected textual data, the corpus of which consists of over a thousand pages of lab research papers (drafts and published work), grant proposals, website material, memos, email exchanges, instant messaging (IM) chats, and internal Wiki board posts.

The data presented here focus on in-situ moments in which the basic–applied boundary was raised during observations, interviews, or in text. I was much more interested in the ways the boundary got deployed in everyday sensemaking than in artificially imposing my own categorization scheme, and so here I emphasize moments in which the boundary arose spontaneously, rather than in response to a direct line of questions. This also means, however, that I present a fairly broad range of evaluative talk about science, technology, and AI. Understandings of the basic–applied boundary were very often mingled with other loosely connected distinctions, such as theory and use, academic and industry/societal, or science and technology. I have not stripped away these interconnections in an attempt to untangle these untidy distinctions. Terminology like 'real AI' and 'real science' was often used interchangeably with basic and applied, as noted below. I highlight these moments as examples of ontological groundings. Since I was tuned to this theme throughout the data collection, I frequently pre-coded data for explicit or implicit definitions of what AI science is or ought to be. After transcription and writing out notes, I used two master codes: scientific demarcation (what is or is not science) and scientific identity (discussion of lab or individual identity), both broadly construed. I then derived inductive sub-codes that included basic/applied, real science, real AI, good/bad science, theory, empirical, data, application, users, comparisons to other labs, and internal/external.

The Deep Reasoning Group (hereafter, DRG)[4] is a lab working on computers that reason with analogy. Its research is based within the AI subfields of knowledge representation and qualitative reasoning. The group consists of a faculty lab head, Derek, two full-time research assistants (one with a PhD and the other an MA), and between 15 and 20 doctoral students. The DRG received the majority of its funding from federal government agencies such as the Defense Advanced Research Projects Agency (DARPA), the Office of Naval Research (ONR), the National Science Foundation (NSF), and the US Department of Homeland Security (DHS). Derek maintained thick ties to research scientists working at these agencies, other researchers in the fields of cognitive psychology and AI, and with military commanders.

The Clever Minds Lab (hereafter, CML) focused on AI for smart information search, retrieval, and processing. The lab consists of two faculty heads, Charles and Cliff, and 10–15 doctoral graduate students. The lab heads were trained primarily in the AI subfield of case-based reasoning (CBR), which is among the most commercially oriented areas of contemporary AI. CBR was a kind of successor to expert systems technologies that had, by the late 1980s, largely fallen out of favor in academia and industry (more on this history below). The CML maintains relations with communication technology firms, consumer product and advertising firms, the arts and entertainment industry, venture capitalists, and city government.

Locating 'real science' in academic AI

It was crucially important to members of the DRG and CML to produce 'real science'. Interestingly, each lab's ontological grounding for their style of AI invalidates their neighbor's, yet on they go enacting their separate realities just steps away from one another. In this section I characterize their contrasting positions, both of which hinge on different renderings of 'real AI'. I will begin, however, with some institutional history.

In the 1980s, expert system technologies were at the center of the most volatile boom, bust, and gradual recovery cycles in the history of AI (see Collins, 1990; Crevier, 1993: 197–216; Edwards, 1996: 294–295). Crevier (1993) notes that expert systems ushered in the first decade in which 'AI stopped being an academic curiosity' (p. 197). Well-known and widely reported examples of expert systems included technologies for medical diagnosis, large-scale industrial production, meteorology, and even the geological location of molybdenum, an ingredient in steel alloys. By the early 1980s, the production of these systems had solidified into a typical form. An AI researcher or team would either solicit or intuit the working knowledge of professional experts and then attempt to program that expertise into a 'knowledge base' of elaborate decision trees. The resulting system, it was hoped, could then assist or possibly replace the human expert. Expert systems attracted considerable industry attention based on their promise of routinizing heretofore non-routine tasks, therefore lowering personnel and training costs. Large corporations such as Boeing, Campbell's Soup, General Electric, and many others invested heavily. Some companies, such as XEROX, founded internal AI groups. Following in the path forged by entrepreneurial scientists in the biomedical sciences

(Shapin, 2008: 222–229), a number of leading AI scientists founded their own private ventures. This included high-profile faculty at leading AI labs such as Carnegie Mellon, MIT, and Yale.

Expert systems, however, largely failed on their promise to scale-up. By the late 1980s, many of these systems exhibited significant bugs, blind spots, and other problems. Critics gained attention in the academic, industrial, and popular press, many suggesting that expert systems could only automate the most machine-like aspects of an expert's task. Expert systems struggled to capture experts' tacit and tactile knowledge as well as with problems with the regression of codified rules (see Collins, 1990: 78–105). By the early 1990s, industry investment in expert systems had shrunk dramatically. Historians refer to the subsequent period as an 'AI Winter'. It is important to note that expert systems were never unanimously embraced within academic AI. A number of prominent AI scientists considered them to be a distraction from the more serious business of understanding intelligence from a more experimental and cognitivist idiom. Nevertheless, the amount of industry money available for their proliferation maintained, for a time at least, their high profile within the field.

A key area of AI to thaw from this winter was CBR, which involves the development of general heuristics, analogies, and representations across similar tasks that can be used toward the accomplishment of a current task. Many of the common application domains of CBR research (e.g. industrial production, international diplomacy, fine cooking) are similar to those used in expert systems. However, CBR techniques are considerably less intensive on the knowledge entry side. They rely on an overall 'script' of a task that can change details as necessary, rather than codified expertise (the formative reports on this reconceptualization are Minsky, 1974; Schank and Abelson, 1975). As such, they tend to be cheaper and easier to build than expert systems since researchers with superficial knowledge can intuit much of the overall task script. Although CBR never produced the same level of industry excitement as expert systems, in part due to more modest goals, CBR researchers have parlayed some of the industry connections established during the height of expert systems AI. There remains, however, an ongoing debate within academic AI over whether CBR systems produce scientific knowledge at all.

DRG staff researcher Donald's desire to create 'real AI' is indicative of this relegation of both expert systems and CBR to technical gadgetry. After completing his dissertation in the early 1990s, Donald began his career at an industry lab that contracts with various companies to build task-specific expert systems and CBR. Initially he enjoyed the exposure to new task environments. During an interview, Donald noted that the work involved 'tangible, real world results'. However, 'it was a good experience but it wasn't research. That's really where my heart is and what I wanted to be doing'. The industry research imposed narrow conceptual constraints and relied on well-worn techniques:

> This was really not AI. It was AI-lite [LAUGHS] … this was the nineties and what happened is the economy started to boom so there was a lot of money for doing things that weren't really AI. It sucked a lot of people away, including me.

In addition, Donald recalled how an instructional team at a contracting company once took credit for his research team's design work when promoting a tool to their clients, which seemed unfair. Donald's eventual move back to academia owed much to his conviction that university research has a more meritocratic reward process for those who develop good ideas. It is also, for Donald, where most of the real AI science gets done.

In our interviews and casual conversations, Donald never set out a clear or settled definition of what he meant by either 'basic science' or 'real AI'. He used these terms interchangeably but both were in contrast to his research experience within industry. He found the latter to be too narrowly focused on profitability and/or utility, not how intelligence operates. Donald's first move back to what he referred to as 'real science' was to join a DARPA-funded project: 'I jumped at [a DARPA project] because it was a chance to get back more into real AI … I worked on a grant for Mixed Initiative Planning for the Air Force'. However, new versions of the basic and applied boundary kept surfacing at DARPA, which Donald felt shifted with the 'revolving door of personnel' at the funding agency. Donald struggled to align his research with the funding agency's constantly moving line between 'out-of-the-box' research and 'field deliverables':

> There was a lot of interest in the military but DARPA was waffling on how applied they wanted to be. DARPA has this problem. They continue to go back and forth between being the civilian research agency and the defense with a capital D research agency. On the one hand, they want things that'll be really easy to sell to the military because that's how they're getting their funding. On the other hand they want things that are out of the box and wild. Blue sky stuff. Trying to find the balance between those two is very hard.

Donald felt, and remains convinced, that his Air Force project was discovery-based research and was therefore quite 'basic'. He was unable to convince DARPA reviewers, however, and it was defunded for being too applied. Nevertheless, a practical effect was to convince his old friend, DRG lab head Derek, to hire him as a staff researcher. During a DRG lab meeting, Donald extolled the virtues of his new academic lab by comparing it directly to their neighbor, the CML: 'We are the kind of place that does science rather than build applications' (field notes).[5]

For Donald, real AI involves the development of frontier techniques rather than a tight execution of well-understood procedures. Producing real science had to do with the knowledge produced and its social organization. That is, real AI research is linked to an economy of authority that attributes credit to the design of a path-breaking technology rather than a technology's immediate utility. Credit should also go to the design team, not individuals conveniently positioned within buyer–seller relationship. Note, however, that Donald's desire to return to what he considers basic science involved a series of nested boundaries. He did not view the systems he built for an industry lab as 'real AI' and preferred the work being funded by DARPA. However, DARPA 'pulled the plug' on his Mixed Initiative project based on a different version of the boundary. Whether or not Donald's project was basic or applied in some absolute sense is largely beside the point. He was unable to produce a research project that DARPA was willing to place into their basic research category. At the same time, the boundary cannot be reduced to mere illusion, propaganda, or a figment of imagination. The 'ontic activity' around Donald's research had material and symbolic consequences. 'Real AI' and its opposite, in this case Donald's Air Force project and his industry work, respectively, were situated within a historical context and a particular configuration of contemporary institutional entanglements.

In contrast to Donald's position, the head of the CML, Charles, considered CBR an entirely appropriate academic successor to the inroads made by expert system's foray into industry relevance. Charles worried that academic AI became what he referred to as 'hallucinatory' once it was removed from 'real-world' contexts of technology use. Charles and CML research were animated by the idea that university–industry entanglements improve both the impact and fundamental insights of AI science. Here, the main impediment to doing 'real AI' was not the narrow utilitarianism or profit motive of industry, but the insular pursuit of canonical theory that has little relevance beyond academic specialists. Charles frequently complained that if an AI system is built for usages beyond its academic home, it is too often dismissed as applied and therefore 'unscientific'. He considered it his main professional goal to develop alternative forms of scientific practice, design, and social organization that reorganize and perhaps even explode what he saw as the imbalance favoring basic over applied knowledge in university research.

In sum, then, Donald and Charles, working just a few doors from one another, configure an AI science that cancels out the reality of their neighbor's work. In a single CS department, we find a proliferation of real AI. Donald uses a distinction between basic and applied that echoes the broad outlines of Vannevar Bush's science policy. Charles has a more problematic relationship with this traditional boundary. In his view, the boundary between basic and applied is not so much illusory as counterproductive. He sees 'real AI' in technologies that unite societal impact with scientific insight. Neither version of 'real AI' is determinative. There is no stable, universally identifiable referent to adjudicate these contrasting ontological claims. Instead, the boundary is positioned and repositioned, assumed and problematized, rendered and re-rendered, as a matter of course. In the sections that follow, I describe the three repertoires that produced basic and applied knowledge at these two labs.

Partitioning

A partitioning repertoire clearly separates basic from applied knowledge according to some mutually exclusive tendency within each. Basic research is rendered, by definition, as different from applied research, and science, in this conception, should be organized to reflect their distinctive features. This repertoire makes basic science sequentially prior to, and/or more fundamental for scientific understanding than, applied science. This is the repertoire, then, that undergirds the substantive position within a 'strong form' of the linear model of innovation (Balconi et al., 2010: 5; Godin, 2006), although this too takes a wide array of forms across settings and projects. The repertoire usually involves an implicit line separating knowledge-for-its-own-sake and knowledge-for-use, although not necessarily. The key feature is that the entities are defined in juxtaposition. For example, if basic research is focused on discovering how a poorly understood entity is organized, then applied research will involve putting the entity to work within an application. Or, if basic science is unfettered by concerns for profitability, then applied science focuses on commercializing discoveries. This provides enough of a definition to move ahead empirically, since I am not trying to deduce these repertoires in the abstract but rather to induce the patterns within which the basic–applied boundary gets orchestrated in rhetoric and practice.

Partitioning was the primary repertoire deployed at the DRG. The DRG placed clear priority on doing research that resulted in scholarly publication, which was seen as the primary indication of scientific knowledge production. Building a technology that could be applied outside the lab was a legitimate goal, but decidedly secondary. Basic research could be conducted without creating an application or usable technology. The DRG's partitioning repertoire was typically conveyed in an informal fashion. It was something that novices within the lab were taught and that required periodic reinforcement. This involved three main lessons. First, the repertoire hinged on the development of an intersubjective understanding that basic science is qualitatively distinct from knowledge-for-use (note that what basic science 'is' at the DRG was never defined and was, in fact, of little relevance). Second, basic science was to be kept methodologically separate from application and with priority accorded to the former. Third, proper sequencing of discovery, design, implementation, and measurement was the key to a successfully executed partitioning repertoire. In the examples below, I first discuss a DRG project that successfully partitioned basic from applied research and then a project that was not successful. These achievements are based on local sensemaking, not a scorecard or other objectified measure of success. Each project yielded academic publications, presentations, and the like. However, DRG members discussed the first as a routine culmination of their work, and the second was treated as a cautionary tale of insufficient partitioning. Last, I discuss the main way that the partitioning repertoire played out at the CML.

The DRG's head, Derek, frequently stressed the importance of separating, prioritizing, and timing the stages of project development. Consider how he aligned the members on his Cognitive Associates project[6] around the staging of performance 'metrics'. Derek sent out the following memo just before a project meeting:

> DARPA wants metrics, not adjectives. And, in fact, we want numbers too – that's the sort of hard-fact details that leads to publishable papers, something we haven't been doing enough of. The subjective nature of many DARPA evaluations leads to less plausible numbers … bringing in experts to do the evaluation is a time-consuming and expensive process. I believe we can get a reasonably objective scoring system in place that will provide a more robust scheme than usually found. (email)

Derek is confident that DARPA's demand for 'hard-fact details' can coexist with his lab's desire to produce scientific publications. However, he emphasized the need to limit (not eliminate) external intrusions on their lab work. He worried that if DARPA brought in outside experts,[7] they might muddy their evaluation of the AI system's performance with subjective impressions of his lab member's performance during the demonstration. Derek wanted to preempt this uncertainty by developing a plan to measure system performance that DARPA representatives would right away recognize as valid. Derek called for a 'reasonably objective scoring system' – being too pragmatic to call for an unassailable one. At the project meeting that followed, Derek went over the goals of Cognitive Associates. He stressed that pleasing funders and doing good research are distinct activities:

> We will not divert to doing demo-specific code. Okay? Those are evil activities … it destroys your soul in the long run. [GROUP LAUGHTER] And in the short run it leaves you with a pile of bailing wire and bullshit code.

Despite this skeptical attitude toward pleasing outsiders, Derek noted that since 'nobody has done something like this, [the DRG, university, users, and DARPA] … will all benefit' (field note).

Parker and Crona (2012) suggest that such imbroglios try to be 'all things to all people', but rarely do they do so all at the same time. Derek repeatedly emphasized the timing and staging of research practice: 'We won't talk about metrics right now because metrics is something you do once you have the system up and running. We are not running yet so we won't measure it' (field note). A proper scientific problem, according to Derek, was one that did not have a readily available answer. Measuring performance too early could ensure failure by limiting the scope of the research problem. Derek simultaneously worried that measuring system performance too late might create the appearance that the team had 'baked in' positive results (interview). That is, for an AI system to be credible it needs proof that it performed an intelligent act, not simply executed a mechanical rule. The longer one waits to test the system's performance, the more opportunities there will be to 'program in' the answers.

Less than a year later, Derek decided the time was ripe to test the system's capabilities. The scheme involved entering graphical representations or 'glyphs' of questions from the Bennett Mechanical Comprehension Test (BMCT), a standardized aptitude test for human engineers. The procedure involved entering known problems and solutions into a knowledge base, commonly referred to as a 'KB', both as glyphs and as formal queries, as shown in Figure 1 from a DRG publication.

The Associate system relied on the DRG's core analogy algorithm to search its knowledge base for a known solution to a problem that involves three specified 'base objects': a person, a rope, and a pole. Once retrieved, the known solution is mapped to a 'framing analogy' by the system's 'reasoner'. Armed with this analogy, it offers the 'suggestion' that 'pole length predicts balance'. The system churns through the following deductions to accomplish this:

> Cognitive Associates makes an analogy from Acrobat A (the base) to Acrobat B (the target) … The system compares Acrobat A's Pole-Length with Acrobat B's Pole-Length, given by the framing analogy … the system answers that the stability has increased from A to B. The system responds to the query by returning the correct answer: (solveDQProblem BMCT-S-17-MEK Object-1297 Object-1298 ((QPQuantityFn Stability) Object-1297) IncreasedDQ) (paraphrased DRG paper)

In practice, this involved 'training'[8] an Associate on solutions to questions likely to appear on a BMCT test. Once the knowledge base included several solutions, Derek told his lab he would use a 'surprise test, Form T, that … [he hadn't] looked at. It … [was] in a drawer in … [his] office'. Form T was a subset of BMCT questions that no one had seen (not even Derek). These tersely communicated details about Form T created considerable intrigue and anxiety among DRG members.

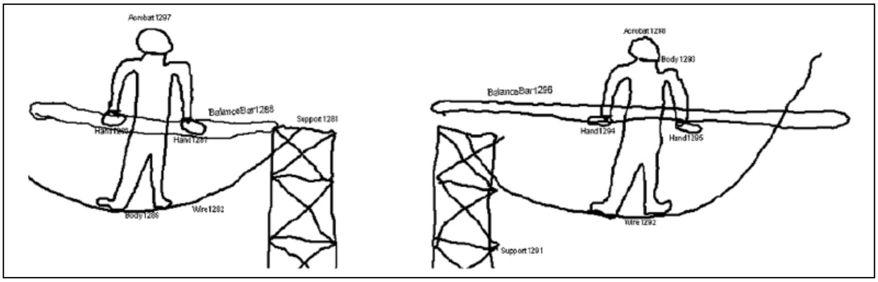

Figure 1. A DRG Analogy Test.
(This Figure) shows a problem from the BMCT sketched in our system. The question, 'Which acrobat will find it easier to keep her balance? (If equal, mark C)', was translated into the following query: (solveDQProblem BMCT-S-17-MEK Object-1297 Object-1298 ((QPQuantityFn Stability) Object- 1297) ?value)).

In terms of its actual performance, the 'experimental results' of the test were decidedly mixed. The system took the Form T exam several times with several different individuals entering glyphs. It was only with the lab members who had considerable familiarity with the glyph entry system that the Cognitive Associates system produced correct answers on 70 percent of the test questions. (Although this was never discussed at the lab, it suggests that a key predictor of success was how well the technical interface had 'trained' the human knowledge enterer). These seemingly modest results did provide data good enough for several conference and journal publications. Moreover, the Associates system and its results were considered a success at the DRG. Results may have been limited, but research goals were appropriately prioritized, the methodology enabled the team to write up the results in an experimentalist idiom, and the papers were published in well-known AI and Cognitive Science venues that DRG members considered purveyors of real AI. Perhaps just as importantly, the modest results enabled the team to express confidence in the system's promise: 'Our goal is to expand [Cognitive Associate's] capabilities via instruction to include the entire BMCT. It would be ground-breaking in qualitative and common sense reasoning to perform at an expert level on such an exam' (paraphrased DRG publication).

Despite Derek's optimism, ferreting out knowledge production from utility was a routine challenge, and not every DRG project was considered a success. Consider Tutor Model, an NSF-funded project that simulated a tutor for formal modeling of causal relationships. Deborah, a DRG graduate student, had primary responsibility for the system's design and implementation. In a conference proceeding, she described the system's goal: 'thinking through causal links within a scenario (e.g., the microwave heats. Heating makes the water's heat increase, which causes its temperature to rise) is a central task for [middle school students]' (paraphrased). The design of Tutor Model involved an attempt to teach middle-school students how to re-represent such causal chains in formal language. At a presentation on her progress, staff researcher Donald expressed significant worries about the project's configuration:

> Donald scratches his head and says, 'There are a lot of confounds here. Who is the target audience for this research? Is it meant to be a teaching tool? Or a way to enhance student learning? Or is it meant to address a set of questions that researchers in education or design have? The questions about the scaffolding seem confounded with all this. Is it just about design?' After a short pause, lab head Derek responds, 'Well, the teachers are not the objects'. Donald suggests that designing an interface is distinct from teaching a substantive domain that middle-school students know little about. Derek is getting impatient with Donald's line of questioning: 'Well, guess what, it turns out they are not. You have to do both'. Donald asks, 'But what have the students learned here? If you take the software away from them, can they write the sentences in the correct way?' Instead of responding to this question, Derek jokes that at least Deborah did not make a 'happy video' with children pretending to have fun. Deborah demurs, 'Your criticisms are legitimate Donald. It's true that a lot of different agendas have gone into this project'. (field note)

Derek blamed the difficulties of Tutor Model on two issues related to partitioning of the basic and applied boundary. First, there were a series of limitations imposed by the Java script used to program the user interface. Derek felt that the interface was a superficial aspect of the technology design, but it was one that was compromising their scientific progress. Second, the system had been implemented at the 'lower end schools' and the 'rattier schools' in the public school system. Derek pointed out that these problems 'had nothing to do with the intellectual interests of our lab'. In short, the context of use had become confused with the lab's ability to produce universal knowledge. Tutor Model was not only an example of basic researchers wishing to ignore the messiness of implementation. It was also an example of the priorities and problems of the implementation environment getting 'confounded' with the priorities and problems of DRG science. That is, it was a DRG project that had insufficiently partitioned basic from applied emphases.

After the meeting, Deborah told me that she appreciated Donald's criticisms but found it difficult to talk about them with the head of the DRG, Derek, present. Deborah had long felt that the system would be compromised unless they could convince the teachers and students to recognize its benefits to them. Deborah felt that by blaming the Java script and dismissing the internal dynamics of the classroom, Derek could avoid seriously interrogating the assumptions of his research program. Deborah noted in an interview that

> The intellectual issues in the way [Derek] sees them, and the way I thought about them originally, is modeling for reuse and transfer. It is a beautiful idea - that if you give somebody a general purpose modeling formalism that is adaptable to multiple situations, are they more likely to recognize an analogy? Are they more likely to make big models out of small models? Can we design this toolkit so that they can use modeling for reuse? (interview)

Deborah had little success getting Derek to see that schools with large numbers of poor and working class minority students might actually strengthen the science of Tutor Model. After all, she reasoned, this might represent a 'hard case' instead of a 'distraction'. Deborah lamented that '[for Derek] the reality of the classroom is not a real thing'. What is real for Derek, she believed, are the universal correlates of human-level intelligence

– how *any* thinking agent uses analogy. In this way, Derek's conception of 'real AI' was too basic for this NSF grant, with basic defined here in terms of universalism. The NSF was too applied, seeking technologies that could measurably improve student learning in specific settings.

Tutor Model was troublesome for the DRG not because it involved a complicated implementation environment, per se, but because its methodological design blurred together too many distal realities: an academic AI lab focused on universal knowledge, resource-deprived public school classrooms, poor neighborhoods with little computing infrastructure and poorly prepared students, well-intentioned but harried teaching staff, and a government agency looking for clear measurements of student learning. In a casual conversation, Deborah pointed out that 'a lot of [the DRG's] systems have been applied in places like engineering schools, where it is much more about content delivery. This is a far more ambitious project'. Deborah finished her dissertation as a collection of separate papers. Several were published in AI and educational research conference proceedings and one in the flagship magazine of the field. Despite these objective successes, Tutor Model was abandoned and its main impact on subsequent DRG projects was as a cautionary tale about the perils of insufficient partitioning.

Unlike the preceding examples from the DRG, the CML usually deconstructed the traditional distinction between basic and applied knowledge. However, the CML did use a partitioning repertoire to manage some of the most financially and legally risky issues of academic capitalism. This was the role of TransferLab, an on-campus incubator for computer systems identified as possessing market potential. Consider how this interstitial unit was used for the development of a CML system named Follow Up. CML graduate student Curt had developed Follow Up as a recommendation system for a university library. In publications, CML members referred to it as an 'automated library research assistant'. Follow Up compiles the information available to the library computer system when a user checks out a book, such as location, item history, and user history, to automatically compile additional resources relevant to the (presumed) user's task. In principle, programmers could build 'wrappers' for Follow Up to provide a similar function for any number of other location-based services. CML head Charles frequently highlighted Follow Up as a CML system with strong commercial potential. The following transpired at a demonstration of Follow Up to representatives of a telecommunications firm:

> One of the visitors asks what it would entail if he wanted to use Follow Up at her firm. Lab head Charles replies that their lab is 'in the midst of licensing this technology for a start-up'. John, a university administrator escorting the guests, asks Charles, 'This technology hasn't officially moved to the TransferLab, right?' Charles replies, 'Not yet, but soon'. The visitor is confused. John explains, 'The TransferLab is an internal incubator within the University. It provides a new model for doing academic research'. Another visitor asks if the TransferLab has for-profit status. Charles replies that it does not since it is part of the university. Charles states that when academic technology gets commercialized it is sometimes cut out of its research phase too quickly. The TransferLab provides a place where 'transfer' can occur without removing the project from the 'discovery phase' too soon. This way, adds Charles, 'Curt doesn't have to stop being a graduate student and I don't have to stop being a professor. It's a powerful model for this kind of research'. (field note)

Boundary organizations like the TransferLab cleave commercial concerns from scientific ones by moving them to a separate unit within the university ecology, similar to how an industrial firm might separate marketing from production. Here, again, the physical separation of space proved salient. The TransferLab involved a formally separate staff working in a separate location from the CML (at a 'research park' adjacent to campus). The TransferLab employed a full-time CEO, an administrative assistant, a Director of Research, two research assistants, and several independent programmers. Prototypes had separate development budgets from the lab funds used in their initial development at the CML. If a technology was thought to be patentable, TransferLab staff handled the patent procedures. If a technology looked as though it could move to market quickly, the TransferLab handled the marketing and product design phases. This separation served symbolic, legal, and financial rationales as much as personnel ones, since both RAs were graduate students in the CML (one was Curt) and the Director of Research was Charles. However, their time and work at the TransferLab were formally independent from their role within the lab and the Computer Science Department. Rather than allow the methodological, legal, and financial pressures of his dissertation research and commercial product development to completely blur, the TransferLab provided a 'trading zone' (Galison, 1996) for managing tensions between scientific knowledge production and commercialization. Boundary organizations like this buffer doctoral students like Curt from direct pressures to produce an immediately saleable product. They are an organizational instantiation of the partitioning repertoire.

Flipping

The CML employed two additional repertoires for rendering basic and applied knowledge that were not in circulation at the DRG. Both problematize the boundary between basic and applied but in different ways. The first, *flipping*, hinged on an inversion of the traditional hierarchy of basic science as knowledge-for-its-own sake and applied science as knowledge-for-use. 'Real science' is grounded in the latter, no longer the former. The flipping repertoire does more than simply state a preference for one style of research over another (although it does this too). It is more thoroughgoing by calling for a broad rethinking of what ought to count as good science and how to reorganize scientific practice to reflect that rethinking. Flipping reverses the sequential and causal assumptions of the linear model of innovation. Knowledge-for-use becomes prior and more important. In this way, the flipping repertoire involves the heterogeneous linkages needed to create what Nowotny (1993) famously referred to as 'socially distributed knowledge'. Canonical theory, disciplinary problem choices, and academic freedom to explore are re-conceived as insular, parochial, out-of-touch, and impotent. Scientific knowledge produced in contexts of use, in turn, is outward-looking, cosmopolitan, current, and impactful.

Both Charles and his lab co-director, Cliff, spent large portions of their days maintaining relations with corporate firms such as Accenture, Diamond Cluster International, General Motors, Honeywell, Northrup Grumman, and SBC Ameritech. Both engaged in consulting for these firms, and Charles even served as a technology advisor for the mayor's office of the large city in which university is located. In addition to these industry and government ties, both men cultivated links with venture capitalists and a vibrant regional arts and entertainment industry. Neither Charles nor Cliff expressed concerns that their scientific practice or the knowledge they produced was getting coopted in this knowledge capitalization process. Instead, they tried to evince the synergistic energy found in the 'About' tab of their lab's website: 'We are always in flux. Our members are the basis of an atmosphere where students, guests and business partners flow through our halls. Together, we create an exciting culture of innovation and impact' (paraphrased). With the themes of flux, flow, and excitement, the CML sought to distance itself from the image of a sterile lab environment. Faculty and staff provide a foundation; 'industry partners', among others, do not distract from knowledge production but enable it. Here, what Pickering (1984) described as 'opportunism in context' is deployed as an explicit strategy for inverting the implicit hierarchy in the traditional basic–applied boundary.

As mentioned earlier, CML members were highly critical of what several members referred to as 'traditional AI'. Their criticisms could be snide and rhetorically savvy, but more importantly they suggested that contexts of use provide a superior launching point for AI science than canonical theory. Consider the following discussion between lab head Charles, co-director Cliff, and a few graduate students about a recent conference they had attended. The topic shifts from a discussion of various lab demonstrations to a more general critique of obscurantism in intellectual life:

> Charles comments, 'One of the funniest things that has happened to us is that we'll get responses to proposals that are like, "This is really applied. You just want to apply stuff you've already done," and so forth. But our proposals are way out of control conceptual. But then we say we think we can do it and base it on this actual technology that we built. It just irritates the hell out of people. I've been thinking about our history of proposals over the last four years. Great proposals. Filled with high concepts and sort of strange, crazy places to go and we've also combined that with this sense of certainty that we can do it. And so people are like, "No, no we want high risk."' Cliff suggests, 'I think that what people confuse is the risk mitigation and the research strategy versus whether the actual end result is a risky thing' and then adds, sarcastically, 'We should throw in more equations'. Laughing, graduate student Cyrus offers, 'Put the word regression in there'. Charles then summarized, 'When we talk to other people in totally different fields it turns out that this is an actual, honest to God ubiquitous problem in the intellectual life of this country. It really is that being unintelligible is synonymous with doing good work'.

Charles, Cliff, and Cyrus are not merely complaining about the rhetoric of basic and applied science in AI (although clearly they are doing this). Their criticism runs deeper. Charles attempts to articulate what is singular about CML research practice. He notes that too often, academic research confuses 'high concepts' and exploration with 'doing good work'. His criticism suggests that CML research practice involves the creation of 'actual technology' that will reliably work. This requires the active building and maintaining of linkages with interested 'partners' outside of the academy. If the technology is 'brittle' (i.e. only functions properly in a tightly constrained environment), Charles points out, then it is very likely that the concepts it is based upon are flawed. Their discussion suggests that the traditional hierarchy between basic and applied science in AI, as well as in 'totally different fields', renders a bias toward substantively empty concepts with little explanatory power and no practical impact. This is a call to reorganize scientific practice in word and deed.

Despite CML member's opposition to what they describe as 'mainstream AI', there is an institutional basis for this style of research within academic AI. For example, taking stock of disciplinary progress over the last 40 years, an editorial survey in the flagship journal *Artificial Intelligence* noted that many scientific advances emerged from 'real-world' applications: 'These are not systems working alone – they must take into account the people with whom they are working – modifying their behavior to fit human constraints and goals, as well as the complexities of the physical world' (Bobrow and Brady, 1998: 1–2). Similarly, the CML's flipping repertoire echoes the critique of top-down computer design launched in the field of informatics. Here, good computer design is measured by the extent to which a system meets user requirements (Jelsma, 2003; Oudshoorn and Pinch, 2003). Advocates of this design philosophy position themselves as political representatives of users by deploying a 'liberal, humanist, and antitechnicist rhetoric' (Cooper and Bowers, 1995: 51) and have developed a number of conferences and publications that reflect this orientation. Echoing a theme prevalent within 'ubiquitous computing', CML members argued that the optimal state of the machine–human relationship is to 'reduce friction' and bring users information 'just in time'. Charles liked to toy with the unease some CML visitors might have with the surveillance potential of proactive search by repeating a lab mantra, 'No matter where you are, no matter what you're doing, no matter what you're thinking, we will get you the information you need … and it turns out that we can figure out what you're thinking' (field note).

Science labs commonly calibrate their research practice and standards of success with those they consider peers (Knorr Cetina, 1999; Traweek, 1988). The nature of these comparisons differed between the DRG and the CML in a revealing way. At the DRG, comparisons focused on how the lab's research on analogical reasoning fits into a division of labor among other labs in AI and the experimental cognitive sciences. CML members, in contrast, considered the combination of their methodological approach and unique *esprit de corps* as the source of their innovative research. That is, CML members focused not on their place among peer labs but on their superiority to them. The flipping repertoire contributed to this spirit of competition. Consider how lab head Charles described the typical CML research process:

> 'Let's build the thing. Let's make it work!'… Rather than, 'Let's hallucinate what the answer is and … then build scaffolding to make sure that is the answer'. We're totally agnostic on how we make things work. We'll just make them work. It's an empirical science … We're letting the practice precede the theory because we think that that's the way computer science should be done … For practice not to precede theory … is like ignoring the physical world in the early days of physics. (interview)

Several members of the CML shared similar sentiments with me. Too often, they reasoned, AI science begins with an answer and then 'bootstraps' the problem rather than starting with a 'real-world' problem. Just as a staff researcher at the DRG, Donald, compared his lab's science favorably to the CML's, members of the CML often characterized their science in contrast to that of the DRG. Charles considered the DRG an example of an outmoded tradition in AI that, like expert systems in the 1980s, relies on massive knowledge elicitation to create uninspiring results. He stated casually that this might be 'wonky science but there is nothing real in it'. Charles once got a lot of laughter at an informal CML lunch when he boasted, 'Give me a week and I'd have [DRG] systems actually working'. Graduate student Charvik took a similar jab at their next door neighbor: '[The DRG] is in the same tradition as a lot of AI research over the past half century. The government invests big money so a robot can pick up a can of Coke'.

Collapsing

The last repertoire I observed involved collapsing the entire distinction between basic and applied knowledge. This repertoire looks quite a bit like Stokes' (1997) concept of 'Pasteur's Quadrant', in which 'use-inspired basic research' involves both a quest for understanding and considerations of use. There is, however, an important difference. Stokes formulates a stable orientation at the conjunction of two continuous variables. His classificatory scheme therefore lends itself to the tendency that Norbert Elias referred to as *zustandreduktion* or the reduction of process to a static condition (Elias, 1978: 112; Elias, 1998: fn 88, p. 45). A multiple ontology framework can avoid placing knowledge into stable categories. Rather, it watches how knowledge is positioned as such. Reality is in the relation, not the essence of the object.

Capturing a repertoire in analysis that is oriented to ontic collapse is a difficult empirical challenge. This is because the repertoire produces the erasure of the very category one seeks to follow (much like pragmatist metaphysics). However, the activity can be followed if we notice that the collapse is justified with an appeal to an irreducible though ineffable whole.[9] This is best demonstrated through an example. Consider how CML head Charles reflected on the skills he looks for in new recruits:

> We tend to like students who have more of a liberal arts education. Who are grounded in the world as opposed to grounded in computer science. Because the technical skills are important. I mean they are crucial. But being able to look at the world in certain kinds of ways is far more important. That's hard to learn. An orientation around people and what they do and how they do it is what you need. If you don't have that, if you have no instincts in that area whatsoever, what you end up doing is working on technical problems that have little to do with anything. It astounds me. The field is rife with that.

Charles is quick to distance his lab from the stereotypical image of the computer scientists: loners and geeks who feel more comfortable expressing themselves in programming code than within human interaction. He suggests that his members have an irreducibly tacit 'feel' for assessing what people need in their computer-mediated lives. What I find most interesting in Charles's quote is how 'instinct' serves to shut down further explanation (Bateson, 2000 [1972]: 38–58). In doing this, Charles positions CML research at the collapse of several common distinctions in forms of knowledge: the explicit with the tacit, the formal with the informal, the ambiguous with the precise, and basic with applied.

This odd coupling of an antipathy toward universals yet grounding in an ineffable whole could lead to confusion around where to place CML research projects. Consider the demonstrations of Sleuth[10] and Follow Up for Northrup Grumman, a military contractor:

> The Northrup Grumman representative asks if Sleuth is finished. Graduate student Cole [whose dissertation was based on Sleuth] replies that it is never finished in the same way that a lot of traditional AI projects get finished. He states, 'We don't build end-user programs … [until they are] out in the world … For us, the theory comes after the deployment'. Graduate student Curt, who assists Cole on Sleuth, adds, 'We could theorize all day. But really we want to get things started and out there in order to then theorize how it gets used'.

Cole's conviction that 'the theory comes after the deployment' and Curt's comment that 'we could theorize all day' distance the CML from traditional AI research. Research that precedes deployment is cast as easy yet brittle. Sound theorizing for CML members is not decontextualized pontification but contextualized inquiry. Later, another company representative asked, with considerable confusion:

> 'So do you use Artificial Intelligence?' After a short silence, Cole replies, 'That's a loaded question!' Other members of the CML chuckle knowingly at one another, which I reflexively join so as not to feel on the outside of an inside joke. Cole continues, 'We use AI as a toolbox. We use it in the system to go out and see how the system is used in the real world. The theory of our research comes from seeing the system deployed. Through deployment we learn how it gets used and how it can change'.

When collapsing the distinction between theory and practice, CML members did not distinguish between discovery and use. Design, deployment, and inquiry are all collapsed into Cole's conceptualization of the 'theory of our research'. In this sense, an evolving analysis of the precise 'implementation environment' is understood to be integral to the iterative design of an AI system and to the knowledge claims that are made with that design.

Departing from his sometimes sardonic tone, CML lab head Charles discussed the collapsing repertoire he practiced in highly idealized terms. Consider how he contrasted his doctoral training with his current research practice during an interview:

> Cliff and I were both trained in AI in a particular point of view. That was case-based reasoning, or reasoning from memory. Reasoning from experience … We had a religion just like everybody else has their religion, and then at one point we decided we were going to lose religion in order to find God. And God is things that work. And we use all sorts of technology and all sorts of ideas and all sorts of approaches in our work … And that actually meant a lot to our lab. It changed us.

Despite his explicit disavowal of religiosity toward any particular AI technique, Charles evoked an evangelical zeal for eclectic research practice. Like his graduate students Curt and Cole, he sees in a collapsed boundary between basic and applied a more real AI science.

Conclusion

This article demonstrates that the reality of basic and applied knowledge is made within practical action. Treating forms of knowledge in this way suggests that scientific objects are rendered rather than discovered. This does not mean that scientists can make any reality they dream up with a bit of savvy rhetoric. The work that goes into making scientific objects factually real is simultaneously epistemological and political (Mol, 2002). It includes regular re-articulations of epistemic identity, research planning, comparisons to peer labs, 'good enough' metrics of proof and evaluation, attempts to convince interlocutors that your science is worth funding, and, of course, long days and sometimes sleepless nights filled with hacking and debugging. This array of practices, and much more, is mobilized to fortify some objects while weakening others. As Wittgenstein (1969) put it, 'knowledge is in the end based on acknowledgement' (p. 378).

The DRG used a partitioning repertoire to separate basic and applied parts of their science in order to maintain a priority on the former while still building systems that might, some day, be adapted to use. Their claims for producing basic scientific knowledge were grounded in an experimentalist idiom of empirical proof, like passing the BMCT, whereas applications of their research were typically left promissory. The CML deployed a version of a partitioning repertoire too, but mostly deployed two other repertoires aimed at deconstructing the traditional boundary altogether. Sometimes they flipped the boundary to indicate that prioritizing the use-context generates better scientific insights than do the insular practices associated with basic research. Other times, they collapsed the distinction altogether, seemingly borrowing a page from some historians who have suggested the boundary is little more than the antiquated propaganda of policy advisors like Vannevar Bush and John Tyndall.

In one sense, my analysis suggests that these AI scientists are similar to research policy analysts themselves, seeking to identify the limits and empirical basis for the categories of basic and applied knowledge. There is, however, an important difference. The repertoires never came in tightly packaged and carefully scoped definitions of the essential characteristics of scientific knowledge. Such determinism was of little use at either lab. The CML, in fact, deployed all three repertoires despite the fact that each one cancels the reality of the other two out. We might conclude, from a perspectival position, that members of the CML were deluded. Or, with multiple ontology, we can see that the repertoires served as epistemic tools for directing ongoing technoscientific entanglements.

Even in the partitioning repertoire, which was the closest approximation I found to the 'linear model of innovation', I did not observe a sequence of activities that map directly onto this model. Even before a research project was underway, lab head Derek and his members anchored their design within at least a vague imaginary of the potential user base. However, members of the DRG did draw on a hard-edged distinction in which basic knowledge is prior to and more important than applied knowledge. This helped them orchestrate, prioritize, and execute their research plans. In other words, it does not much matter that the 'linear model' is a flawed way to understand actual research practice (Balconi et al., 2010; Edgerton, 2004) since it sets out boundary conditions that are used to affirm epistemic identity and strategize moves within a research field (Hounshell, 2004). Such a model of scientific work is neither fully substantive, in the sense of reflecting a singular reality of how science gets done, nor is it mere artifice, in the sense of being merely a rhetorical flourish or ideological perspective. The model has substantive and artifactual features in its daily enactments. It has a real form of life, but a form tweaked within particular configurations of technoscience rather than turned out as if from a Jell-O mold.

Scientific practice is not unpredictable so much as processual. Some parts of a configuration will have more influence over the basic–applied boundary than others. Lab heads have the greatest strength in determining the parameters of their group's entanglements (which is not to suggest lab heads are *Übermenschen* but that they are made accountable). Furthermore, as legal scholars have pointed out, repeat players more easily impose expectations than do one-shotters in complex systems (e.g. Berrey et al., 2012). The DRG is far more dependent on DARPA for resources than the agency is in need of DRG systems. As such, DARPA is a fairly intransigent force in what counts as basic and applied. DRG staff researcher Donald failed to convince DARPA that his Air Force project was basic enough and so development ceased. Likewise, Deborah's Tutor Model failed to convince senior members of the DRG that she had sufficiently separated the basic and applied parts of her research. The DRG had thin ties of little internal consequence with inner city schools and found that particular implementation environment too messy and distracting. Charles, Cliff, and their CML members struggled, in contrast, to shift what counts as good science from abstract theory to workable systems. They have, however, carved out enough of a niche within their professional field that they can keep this version of AI science going. These two very different realities coexist right next door to one another. They sometimes mingle, particularly when members of the two labs draw direct comparisons to one another. Each version of the boundary cancels the other two out, yet on they go. A fruitful arena for further research is single research projects that combine multiple realities of basic and applied research within them, such as inter-lab collaborations, as well as transdisciplinary projects that cobble together quite different ontic activities.

There is no simple way to adjudicate the question, 'What is basic and applied science?' The answer cannot be deduced by better specifying the substantive or the artifactual elements of the boundary. Doing so, as numerous research policy analysts have with concepts like 'Mode II science', 'translational research', 'socially distributed knowledge', or 'use-inspired basic research', merely carries forth the ongoing activity of partitioning, flipping, or collapsing the basic–applied boundary. Adding new categories provides new objects to be manipulated. However, in this desire to distill, we tend to lose sight of the fact that there is no God to separate, once and for all, the pure from the contaminated, the primary from the secondary, or the basic from the applied. There are, however, devilishly rich answers to this question in the details, way down in the muck of it all.

Acknowledgements

The author thanks the three anonymous reviewers and Sergio Sismondo for their feedback that dramatically improved the quality and clarity of the analysis. The author is especially grateful to Carol Heimer and Art Stinchcombe, both of whom manage to be intellectual mentors, sources of scholarly inspiration, and down-to-earth comrades in one fell swoop. Thanks also to Jorge Arditi, Ellen Berrey, Beth Popp Berman, Alan Czaplicki, Paul Durlak, Mike Farrell, Gary Alan Fine, Daniel Lee Kleinman, Jennifer Light, Jason Owen-Smith, Mike Sauder, and Jared Strohl, all of whom helped at one juncture or another. Finally, thanks to the members of the labs described in this article, who put up with an itinerant observer's pesky questions, incessant note taking, and clunky recording equipment with tolerance and good wit.

Funding

Data collection benefited from the financial support of the MacArthur Foundation, the Mellon Foundation, the Kaplan Center for the Humanities, and the Department of Sociology at Northwestern University. Finishing the write-up depended upon a fellowship from the Humanities Institute and the Office of the Vice President for Research at the University at Buffalo, SUNY.

Notes

1. This does some disservice to the diverse treatments of the basic–applied distinction. For example, research policy analysts like Henry Etzkowitz and his colleagues vacillate between substantive and artifactual positions (not unlike the AI scientists studied below). Abstractions are necessarily simplifications of wholes, and this one is useful for drawing out the ontologies embedded within categorical distinctions of science.
2. The term repertoire comes from Nigel Gilbert and Michael Mulkay's (1982, 1984) work on the discursive strategies that biochemists draw upon to warrant their views of acceptable and unacceptable research findings. A repertoire is a stock of skills, moves, and routines that an actor or a group of actors is prepared to perform within a particular scenario. Gilbert and Mulkay's empirical analysis of repertoires is almost entirely discursive. Mine combines discourse with practice, treating their relation as a joint performance. In a recent commentary on Patrik Aspers's (2014) critique of the ontological turn in Science and Technology Studies (STS), Sergio Sismondo (in press) formulates multiple ontology as examining how 'suites of connected practices establish ways of being … [a] focus on local stability as a temporary narrowing of interests, which sets aside the study of flux and its stabilization precisely to establish the ubiquity of ontic activity'. This is a helpful formulation. My use of repertoire refers to the stock of discursive moves that articulate, direct, and prioritize 'suites of connected practices' for warranting acceptable scientific findings.
3. The concept of a research 'border zone' is based on Peter Galison's (1996, 1997) notion of 'trading zones' but differs in emphasis. I am drawing attention to the sensemaking processes that separate, reorder, or collapse the distinction between forms of knowledge. Galison focuses on how simulation technologies enable cross-domain communication across diverse institutional domains.
4. The names of each lab and their personnel, as well as research projects, are pseudonyms. Quotations from published lab papers have been paraphrased for anonymity, which is noted within the text. To help the reader track members of each lab, I refer to all Deep Reasoning Group (DRG) members with names that begin with D, such as Derek, Donald, and Deborah. I refer to all members of the Clever Minds Lab (CML) with names that begin with the letter C, as in Charles, Cliff, and Charvik. I use first names only, without formal titles, because in daily interaction everyone related on a first name basis, exhibiting the studied informality across rank often characteristic of American academic life. However, to signal status differences, I introduce each individual as 'lab head', 'staff researcher', or 'graduate student' and provide reminders of formal status where relevant.
5. Comparisons like Donald's were common at labs' meetings, in casual conversations, and during interviews, but never occurred in shared spaces. This was one of the key ways that the relative privacy of lab space was instrumental in the production of each group's scientific work even when one version of AI science dismissed the reality of their neighbor's version. The organization of physical space helped orchestrate ontological pluralism.

6. Derek described his Cognitive Associates as building digital analogs to the M*A*S*H character Radar O'Reilly – nerdy sidekicks who pop up with useful information just when they are needed and then go away. The general idea is to create general-purpose digital assistants that would steadily add to a 'learning database' by monitoring how their human partners make decisions.
7. It is customary that in large-scale Defense Advanced Research Projects Agency (DARPA)–funded projects, members of the grant oversight office do periodic 'site visits' to evaluate progress. During my 3 years of observation, the director of DARPA Information Processing Techniques Office (IPTO) visited the DRG twice. On both visits, the director spent a day viewing system demonstrations at the DRG, a second day giving a talk to the Department, and then briefly visited other labs on campus (including the CML). Historically, the Director of DARPA IPTO holds a PhD in the field and is an active Artificial Intelligence (AI) research scientist. This was true of the Director of IPTO in the mid-2000s, who spent much of his career leading an AI research group at a large communications corporation. During his tenure at IPTO, this director's primary initiative was a push for more research on 'Cognitive Systems', which was the stream of funding that supported DRG work. In 2010, IPTO was merged with another DARPA office and became the Information Innovation Office.
8. One might think that 'programming' and 'coding' are the more accurate word choices here. However, DRG members did not use these terms to describe computer capacities that they considered an example of learning or intelligence. Terms like coding and 'hacking' were used to refer to rote procedure. 'Training' was used for human-level intelligence.
9. One might note a parallel here with the actor-network theory's ontological grounding in a 'seamless web' of human and non-human relations.
10. Sleuth involved an attempt to automate information and Internet search by 'pro-actively' filtering a computer user's currently running software applications. It did this by pulling a 'word chunk' or extract from the corpus of words available to it (Word documents, web browsers, recently opened emails, etc.). The chunk excluded 'stop words' such as 'the', 'a', 'is', and so on. Sleuth seeks out and delivers additional web resources, data, news, videos, blogs, and so on, without the user having to come up with search terms on their own. CML members pointed out that search terms are the key source of 'friction' in current search technology like Google.

References

Abbott A (2001) *Chaos of Disciplines*. Chicago, IL: University of Chicago Press.

Alexander JC (1982) *Theoretical Logic in Sociology*. Berkeley, CA: University of California Press.

Aspers P (2014) Performing ontology. *Social Studies of Science*. Epub ahead of print 25 September. DOI: 10.1177/0306312714548610.

Balconi M, Brusoni S and Orsenigo L (2010) In defence of the linear model: An essay. *Research Policy* 39(1): 1–13.

Bateson G (2000 [1972]) *Steps to an Ecology of Mind*. Chicago, IL: University of Chicago Press.

Berrey E, Hoffman SG and Nielsen LB (2012) Situated justice: A contextual analysis of fairness and inequality in employment discrimination litigation. *Law & Society Review* 46(1): 1–36.

Bijker WE (1995) *Of Bicycles, Bakelites, and Bulbs: Toward a Theory of Sociotechnical Change*. Cambridge, MA: MIT Press.

Blumer H (1969) *Symbolic Interactionism: Perspective and Method*. Englewood Cliffs, NJ: Prentice Hall, Inc.

Bobrow DG and Brady JM (1998) Artificial intelligence 40 years later. *Artificial Intelligence* 103(1–2): 1–4.

Bush V (1945) *Science the Endless Frontier: A Report to the President*. Washington, DC: United States Government Printing Office.

Calvert J (2006) What's special about basic research? *Science, Technology & Human Values* 31(2): 199–220.

Collins HM (1985) *Changing Order: Replication and Induction in Scientific Practice*. Chicago, IL: University of Chicago Press.

Collins HM (1990) *Artificial Experts: Social Knowledge and Intelligent Machines*. Cambridge, MA: The MIT Press.

Collins HM and Evans RE (2002) The third wave of science studies: Studies of expertise and experience. *Social Studies of Science* 32(2): 235–296.

Cooper G and Bowers J (1995) Representing the user: Notes on the disciplinary rhetoric of human-computer interaction. In: Thomas PJ (ed.) *The Social and Interactional Dimensions of Human-Computer Interfaces*. Cambridge: Cambridge University Press, pp. 48–66.

Crevier D (1993) *AI: The Tumultuous History of the Search for Artificial Intelligence*. New York: Basic Books.

Dewey J (1922) *Human Nature and Conduct*. New York: Henry Holt and Company.

Edgerton D (2004) 'The linear model' did not exist: Reflections on the history and historiography of science and research in industry in the Twentieth Century. In: Grandin K, Wormbs N and Widmalm S (eds) *The Science-Industry Nexus: History, Policy, Implications*. Sagamore Beach, MA: Science History Publications, pp. 31–57.

Edwards P (1996) *The Closed World: Computers and the Politics of Discourse in Cold War America*. Cambridge, MA: The MIT Press.

Elias N (1978) *What is Sociology?* (trans. S Mennell and G Morrissey). New York: Columbia University Press.

Elias N (1998) *Norbert Elias: On Civilization, Power and Knowledge: Selected Writings* (eds S Mennell and J Goudsblom). Chicago, IL: University of Chicago Press.

Emirbayer M and Mische A (1998) What is agency? *American Journal of Sociology* 103(4): 962–1023.

Etzkowitz H, Webster A, Gebhardt C, Regina B and Terra C (2000) The future of the university and the university of the future: Evolution of ivory tower to entrepreneurial paradigm. *Research Policy* 29(2): 313–330.

Fine GA (1993) Ten lies of ethnography: Moral Dilemmas of Field Research. *Journal of Contemporary Ethnography* 22(3): 267–294.

Forman P (2007) The primacy of science in modernity, of technology in postmodernity, and of ideology in the history of technology. *History and Technology* 23(1): 1–152.

Frank DJ and Gabler J (2006) *Reconstructing the University: Worldwide Shifts in Academia in the 20th Century*. Stanford, CA: Stanford University Press.

Galison P (1996) Computer simulations and the trading zone. In: Galison P and Stump D (eds) *The Disunity of Science: Boundaries, Contexts, and Power*. Stanford, CA: Stanford University Press, pp. 118–157.

Galison P (1997) *Image and Logic: A Material Culture of Microphysics*. Chicago, IL: University of Chicago Press.

Geertz C (1973) Thick description: Toward an interpretive theory of culture. In: Geertz C (ed.) *The Interpretation of Cultures: Selected Essays*. New York: Basic Books, pp. 3–30.

Gibbons M, Limoges C, Nowotny H, Schwartzman S, Scott P and Trow M (1994) *The New Production of Knowledge: The Dynamics of Science and Research in Contemporary Societies*. Thousand Oaks, CA: SAGE.

Gieger RL and Sá CM (2009) *Tapping the Riches of Science: Universities and the Promise of Economic Growth*. Cambridge, MA: Harvard University Press.

Gieryn TF (1983) Boundary-work and the demarcation of science from non-science: Strains and interests in professional ideologies of scientists. *American Sociological Review* 48(6): 781–795.

Gieryn TF (1999) *Cultural Boundaries of Science: Credibility on the Line*. Chicago, IL: University of Chicago Press.

Gilbert GN and Mulkay M (1982) Warranting scientific belief. *Social Studies of Science* 12(3): 383–408.

Gilbert GN and Mulkay M (1984) *Opening Pandora's Box: A Sociological Analysis of Scientists' Discourse*. Cambridge: Cambridge University Press.

Godin B (2006) The linear model of innovation: the historical construction of an analytical framework. *Science, Technology & Human Values* 31(6): 639–667.

Hoffman SG (2011) The new tools of the science trade: Contested knowledge production and the conceptual vocabularies of academic capitalism. *Social Anthropology* 19(4): 439–462.

Hounshell D (2004) Industrial research: Commentary. In: Grandin K, Wormbs N and Widmalm S (eds) *The Science-Industry Nexus: History, Policy, Implications*. Sagamore Beach, MA: Science History Publications, pp. 59–68.

Jelsma J (2003) Innovating for sustainability: Involving users, politics and technology. *Innovation* 16(2): 103–116.

Knorr Cetina K (1999) *Epistemic Cultures: How the Sciences Make Knowledge*. Cambridge, MA: Harvard University Press.

Krimsky S (2003) *Science in the Private Interest: Has the Lure of Profits Corrupted Biomedical Research?* Lanham, MD: Rowman & Littlefield Publishers.

Lam A (2010) From 'ivory tower traditionalists' to 'entrepreneurial scientists'? Academic scientists in fuzzy university-industry boundaries. *Social Studies of Science* 40(2): 307–340.

Lamont M (2009) *How Professors Think: Inside the Curious World of Academic Judgment*. Cambridge, MA: Harvard University Press.

Latour B (1987) *Science in Action*. Princeton, NJ: Princeton University Press.

Latour B (1993 [1991]) *We Have Never Been Modern* (trans. C Porter). Cambridge, MA: Harvard University Press.

Lehrer J (2011) How friends ruin memory: The social conformity effect. *Wired/Condé Nast*. Available at: http://www.wired.com/2011/10/how-friends-ruin-memory-the-social-conformity-effect/ (accessed on 23 December 2014).

Lomnitz LA and Cházaro L (1999) Basic, applied and technological research: Computer science and applied mathematics at the National Autonomous University of Mexico. *Social Studies of Science* 29(1): 113–134.

McCarthy J, Minsky ML, Rochester N and Shannon CE (1955) *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*. Hanover, NH: Dartmouth College, Rockefeller Foundation.

Marcus GE (1986) Afterword: Ethnographic writing and anthropological careers. In: Cliffor J and Marcus GE (eds) *Writing Culture: The Poetics and Politics of Ethnography*. Berkeley, CA: University of California Press, pp. 262–266.

Minsky M (1974) *A framework for representing knowledge*. Massachusetts Institute of Technology Artificial Intelligence Memo No. 306. Available at: http://18.7.29.232/bitstream/handle/1721.1/6089/AIM-306.pdf?sequence=2 (accessed 29 November 2014).

Mol A (2002) *The Body Multiple: Ontology in Medical Practice*. Durham, NC: Duke University Press.

Moore K, Kleinman DL, Hess D and Frickel S (2011) Science and neoliberal globalization: A political sociological approach. *Theory and Society* 40(5): 505–532.

Murray F (2010) The Oncomouse that roared: Hybrid exchange strategies as a source of distinction at the boundary of overlapping institutions. *American Journal of Sociology* 116(2): 341–388.
Nowotny H (1993) Socially distributed knowledge: Five spaces for science to meet the public. *Public Understanding of Science* 2(4): 307–319.
Oudshoorn N and Pinch T (eds) (2003) *How Users Matter: The Co-Construction of Users and Technology*. Cambridge, MA: The MIT Press.
Owen-Smith J (2003) From separate systems to a hybrid order: Accumulative advantage across public and private science at Research One universities. *Research Policy* 32(6): 1081–1104.
Parker J and Crona B (2012) On being all things to all people: Boundary organizations and the contemporary research university. *Social Studies of Science* 42(2): 262–289.
Pickering A (1984) *Constructing Quarks: A Sociological History of Particle Physics*. Chicago, IL: University of Chicago Press.
Pinch TJ and Bijker WE (1987) The social construction of facts and artifacts: Or how the sociology of science and the sociology of technology might benefit each other. In: Bijker WE, Hughes TP and Pinch T (eds) *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology*. Cambridge, MA: The MIT Press, pp. 17–50.
Popp Berman E (2012) *Creating the Market University: How Academic Science Became an Economic Engine*. Princeton, NJ: Princeton University Press.
Rhoten D and Powell WW (2007) The frontiers of intellectual property: Expanded protection versus new models of open science. *Annual Review of Law and Social Science* 3: 345–373.
Rhoten DR and Powell WW (2010) Public research universities: From land grant to Federal grant to patent grant institutions. In: Rhoten D and Calhoun C (eds) *Knowledge Matters: The Public Mission of the Research University*. New York: Columbia University Press, pp. 319–345.
Schank RC and Abelson RP (1975) *Scripts, Plans, and Knowledge*. New Haven, CT: Yale University.
Schütz A (1945) On multiple realities. *Philosophy and Phenomenological Research* 5(4): 533–576.
Sewell WH Jr (1992) A theory of structure: Duality, agency, and transformation. *American Journal of Sociology* 98(1): 1–29.
Shapin S (2008) *The Scientific Life: A Moral History of a Late Modern Vocation*. Chicago, IL and London: University of Chicago Press.
Shore C and McLauchlan L (2012) 'Third mission' activities, commercialisation and academic entrepreneurs. *Social Anthropology* 20(3): 267–286.
Sintonen M (1990) Basic and applied sciences – Can the distinction (still) be drawn? *Science Studies* 3(2): 23–31.
Sismondo S (in press) Ontological turns, turnoffs and roundabouts. *Social Studies of Science*.
Slaughter S and Rhoades G (2004) *Academic Capitalism and the New Economy: Markets, State, and Higher Education*. Baltimore, MD: The Johns Hopkins University Press.
Stinchcombe AL (2001) *When Formality Works: Authority and Abstraction in Law and Organizations*. Chicago, IL: University of Chicago Press.
Stokes DE (1997) *Pasteur's Quadrant: Basic Science and Technological Innovation*. Washington, DC: Brookings Institution Press.
Strathern M (2005) *Partial Connections*. Walnut Creek, CA: Rowman Altamira.
Sturdy S (2007) Knowing cases: Biomedicine in Edinburgh, 1887–1920. *Social Studies of Science* 37(5): 659–689.
Thursby J and Thursby M (2010) University licensing, Harnessing or tarnishing faculty research? *Innovation Policy and the Economy* 10(1): 159–189.
Traweek S (1988) *Beamtimes and Lifetimes: The World of High Energy Physicists*. Cambridge, MA: Harvard University Press.

Vallas S and Cummins E (2014) Relational models of organizational inequalities: Emerging approaches and conceptual dilemmas. *American Behavioral Scientist* 58(2): 228–255.
Vallas SP and Kleinman DL (2008) Contradiction, convergence and the knowledge economy: The confluence of academic and commercial biotechnology. *Socio-Economic Review* 6(2): 283–311.
Van Looy B, Callaert J and Debackere K (2006) Publication and patent behavior of academic researchers: Conflicting, reinforcing or merely co-existing? *Research Policy* 35(4): 596–608.
Whitehead AN (1938) *Modes of Thought*. New York: Macmillan.
Wittgenstein L (1969) *On Certainty* (trans. D Paul and GEM Anscombe). Oxford: Blackwell.

Author biography

Steve G Hoffman received his PhD in sociology at Northwestern University and is currently an Assistant Professor of Sociology at the University at Buffalo, SUNY. His empirical research examines situated decision-making and ontology in a variety of information-rich settings, including academic–industry entanglements, nuclear energy policy, boxing gyms, and small rural towns. He has published in an eclectic array of scholarly venues, but mostly in sociology journals. He does not normally refer to himself in the third person, but does derive modest joy from gently poking fun at this odd form of scholarly legitimation.