

Why National eHealth Programs Need Dead Philosophers: Wittgensteinian Reflections on Policymakers' Reluctance to Learn from History

TRISHA GREENHALGH, JILL RUSSELL,
RICHARD E. ASHCROFT, and WAYNE PARSONS

Queen Mary University of London

Context: Policymakers seeking to introduce expensive national eHealth programs would be advised to study lessons from elsewhere. But these lessons are unclear, partly because a paradigm war (controlled experiment versus interpretive case study) is raging. England's \$20.6 billion National Programme for Information Technology (NPfIT) ran from 2003 to 2010, but its overall success was limited. Although case study evaluations were published, policymakers appeared to overlook many of their recommendations and persisted with some of the NPfIT's most criticized components and implementation methods.

Methods: In this reflective analysis, illustrated by a case fragment from the NPfIT, we apply ideas from Ludwig Wittgenstein's postanalytic philosophy to justify the place of the "n of 1" case study and consider why those in charge of national eHealth programs appear reluctant to learn from such studies.

Findings: National eHealth programs unfold as they do partly because no one fully understands what is going on. They fail when this lack of understanding becomes critical to the programs' mission. Detailed analyses of the fortunes of individual programs, articulated in such a way as to illuminate the contextualized talk and action ("language games") of multiple stakeholders, offer unique and important insights. Such accounts, portrayals rather than models, deliver neither statistical generalization (as with experiments) nor theoretical generalization (as with multisite case comparisons or realist evaluations). But they do provide the facility for *heuristic generalization* (i.e., to achieve a clearer understanding of what is going on), thereby enabling more productive debate about eHealth programs' complex, interdependent social practices. A national

Address correspondence to: Trisha Greenhalgh, Yvonne Carter Building, 58 Turner Street, London E1 2AB (email: p.greenhalgh@qmul.ac.uk).

eHealth program is best conceptualized not as a blueprint and implementation plan for a state-of-the-art technical system but as a series of overlapping, conflicting, and mutually misunderstood language games that combine to produce a situation of ambiguity, paradox, incompleteness, and confusion. But going beyond technical “solutions” and engaging with these language games would clash with the bounded rationality that policymakers typically employ to make their eHealth programs manageable. This may explain their limited and contained response to the nuanced messages of in-depth case study reports.

Conclusion: The complexity of contemporary health care, combined with the multiple stakeholders in large technology initiatives, means that national eHealth programs require considerably more thinking through than has sometimes occurred. We need fewer grand plans and more learning communities. The onus, therefore, is on academics to develop ways of drawing judiciously on the richness of case studies to inform and influence eHealth policy, which necessarily occurs in a simplified decision environment.

Keywords: eHealth, policymaking, case study, ethnography, evaluation, Wittgenstein, sensemaking, learning community.

eHealth Policy: The Canon of History

Those who ignore history are doomed to repeat it.

—George Santayana (1863–1952)

NATIONAL EHEALTH PROGRAMS RARELY UNFOLD AS predicted, especially when carefully planned out in advance. Of course, that is because they are complex and unpredictable. But policymakers often persist in thinking that things will go better next time. Their hubris has reached a level that deserves to be researched in its own right. To that end, this article argues that lessons are rarely learned from national eHealth programs because insufficient value is placed on in-depth case studies, and it makes this case on philosophical rather than methodological grounds. We propose that national eHealth programs and, by extension, other complex technology projects with multiple stakeholders and interdependencies could usefully be reconceptualized as Wittgensteinian language games.

The United States' 2009 Health Information Technology for Economic and Clinical Health (HITECH) Act allocated up to \$27 billion for the adoption and "meaningful use" of electronic health records by physicians and hospitals between 2011 and 2015 (Blumenthal and Tavenner 2010). Australia's federal budget for 2010/2011 included A\$466.7 million (US\$473 million) for the design, building, and national rollout of a personally controlled electronic health record (PCEHR) (Westbrook and Braithwaite 2010). By 2010, Canada's Health Infoway implementation program had been allocated C\$2.13 billion (US\$2.16 billion) (Whitt 2010).

In some ways, England was ahead of the game. Between 1997 and 2010, the United Kingdom Labour government (which in 1948 introduced the National Health Service as part of a cradle-to-grave welfare state) sought to modernize public-sector services with the help of "state-of-the-art" information technology. By this was meant large, centrally procured systems developed by commercial software suppliers working under contract according to detailed advanced specification and stringent technical security standards. The National Programme for IT (NPFIT) in England (though, notably, not in Scotland, Wales, or Northern Ireland) was a paradigm case of such policy (Department of Health 2005). Although it was described by some as world leading in its scope, vision, and technical sophistication, it was dismissed by others as monolithic, inflexible, resource hungry, and overgoverned (Kreps and Richardson 2007). What was not disputed was its substantial cost (£12.7 billion [US\$20.6 billion] over six years) and the fact that its rollout fell progressively behind its widely publicized implementation schedule (Greenhalgh et al. 2010a, 2010c; National Audit Office 2011; Robertson et al. 2010).

In May 2010, a general election in the United Kingdom produced a hung parliament followed by a hastily aligned coalition between the Conservative Party (which has traditionally leaned to the right and sought to roll back the state and to support private enterprise) and the Liberal Democrat Party (which has traditionally leaned to the left and sought to protect civil liberties). Many people expected that these odd bedfellows would soon dismantle the centralized, state-driven NPFIT in favor of smaller, more bespoke systems that would gain in agility what they lost in interoperability and would emphasize local record linkage (e.g., between general practice and nearby hospitals) rather than national integration.

This anticipated shift did not occur, at least not to the extent that many stakeholders hoped. While there was much talk of “decentralization” and “flexibility,” national contracts with commercial suppliers were not canceled (Collins 2010), and two of the NPfIT’s most unpopular technologies—the Summary Care Record (SCR, an extract from a patient’s personal medical record, stored on a national database) and HealthSpace (a personal health organizer that allows an individual to view his or her own Summary Care Record on the Internet)—were retained as central components of the new national eHealth policy that replaced the NPfIT (Department of Health 2010). Conservative and Liberal Democrat politicians, who occupied the opposition benches when the NPfIT emerged and took shape, had, at the time, repeatedly called for the government to be held to account for the program’s high costs and allegedly weak performance. For example, “At a time when every penny of public money needs to be spent wisely, [the prime minister] wants to waste £13 billion on an NHS computer system that does not work” (Nick Clegg, leader, Liberal Democrat Party, Prime Minister’s Questions, October 29, 2008). Yet when Clegg became deputy prime minister in May 2010, he did not pursue this argument and appeared to acquiesce with the opposing position.

The independent evaluation of the Summary Care Record and HealthSpace programs by two of the authors of this article (TG and JR) and other collaborators (in this account, for simplicity, referred to as “we”) followed an in-depth, mixed-method case study design for three years, 2007 to 2010, involving more than 140 interviews, two thousand hours of ethnographic observation, and a statistical analysis of a data set of more than 400,000 consultations (Greenhalgh et al. 2010a, 2010b, 2010c). The evaluation was both formative (providing ongoing feedback to policymakers and implementation teams as the work unfolded) and summative (delivering a final report to which the Department of Health was expected to offer a formal response). That report was submitted, coincidentally, on the day of the United Kingdom general election (May 6, 2010). In the report, we described a host of technical, operational, economic, and political complexities, which, we suggested, explained why adoption of the technologies was orders of magnitude lower than predicted. We proposed that the anticipated “tipping point,” at which everything was going to get easier, was a simplistic illusion. That is, the difficulties we had identified might worsen rather than be resolved. And we reflected on our own ambiguous position as both the

officially sanctioned narrators of this unfinished epic and key characters within it.

Although the English Department of Health, which commissioned and funded our evaluation, formally “welcomed” our report in June 2010, it commented on only two of our recommendations. Specifically, the department chose to focus on the content of the record and the opt-out procedure for dissenting patients, and to that end, it immediately commissioned two further, but much smaller, evaluations. Both were conducted by senior civil servants (1) to define what data fields the Summary Care Record should contain and (2) to review the opt-out process. This left unexplored ten additional areas that our evaluation report (and, in the months leading up to it, our formative feedback to strategy groups within Connecting for Health) had flagged as in need of prompt review, including the sheer scale and complexity of the programs (which, coupled with an inflexible, milestone-driven change model, militated strongly against their ultimate success); the conflicting and often incommensurable perspectives of multiple stakeholders; the questionable wisdom of prominent government involvement; and the numerous tensions and paradoxes, many of which we had classified as “wicked problems,” that is, contested, politically charged, value laden, and inherently insoluble. In October 2010, the English health minister, Simon Burns, announced:

I am pleased that a consensus has emerged about the importance of the SCR in supporting safe patient care, as long as the core information contained in it is restricted to medication, allergies and adverse reactions. Coupled with improvements to communication with patients which reinforce their right to opt out, we believe this draws a line under the controversies that the SCR has generated up to now. (S. Burns, Department of Health press release, October 11, 2010, http://www.dh.gov.uk/en/MediaCentre/Pressreleases/DH_120379)

This statement perhaps reflects a “sociology of expectations” (Brown 2007), that all technological projects, being future oriented, need to create strong expectations in the minds of users, stakeholders, and sponsors in order to enroll and coordinate them. Even though we had deliberately written our evaluation with multiple audiences in mind and intended each to draw on it in different ways, Connecting for Health (the branch of the Department of Health responsible for implementing the NPfIT) put substantial resources into producing a single, coherent, account of the

programs, including releasing tightly coordinated responses to media reports.

The new government's decision not to directly address most of the recommendations in our report surprised us, although in retrospect we should have been aware that previous governments had responded similarly to official reports on the NPfIT (Collins 2007). When the funding was first allocated to the Summary Care Record and HealthSpace programs, a ministerial task force had recommended an academic evaluation, independently commissioned and peer reviewed, as a condition for approving the programs (Cayton 2006). Thus, the Labour government had not only acknowledged our work but also used it as part of the "new transparency" to legitimize its suite of large-scale technology programs (Fisher 2010).

Four possible explanations could account for the new coalition government's decision to focus on the simple and positive messages from its internally commissioned reports rather than on the complex and nuanced conclusions from our independent evaluation. First, senior civil servants may have judged our work to be of poor quality, although it was reassuring that independent peer reviewers had not. Second, the Summary Care Record and HealthSpace may have become "runaway technology projects" at which the sponsor continued to throw good money after bad rather than face up to sunk costs (Sarup 2003). Third, the scale and complexity of the programs may have created a high degree of path dependency (Wilsford 1994). In other words, the programs may have reached a point where it was difficult to do anything other than continue in the direction already begun (e.g., canceling contracts worth more than £500 million [US\$810 million] without incurring prohibitive penalties is no simple matter). Fourth, policymakers who had been steeped in the logic of technological determinism (i.e., that technology X will have impact Y and that Y can be measured) and trained to expect unambiguous, quantitative estimates of "impact" simply did not understand (or, perhaps, did not take time to come to understand) the more nuanced and contingent messages contained in our 234-page case study or even its twenty-two-page executive summary. In this article, we wish to explore this last explanation and suggest that it underlies the first three.

The remainder of this article is the result of reflections by four authors, two of whom (WP and RA) were not involved in the original evaluation, and consists of five main sections. First, we briefly acknowledge positivist critiques of nonexperimental research, mainly to depict the contested knowledge claims in this field. Next, we consider the main

approaches to case study traditionally taken by organizational and social researchers. Third, we outline a justification proposed by Haridimos Tsoukas, who in turn draws on Wittgenstein, for the careful study of a single case to elucidate complex social phenomena, an approach that plays down the theory-building task (analytic generalization) traditionally seen as central to academic case study research. Fourth, we apply this approach to a small fragment of the NPfIT evaluation, so as to illustrate the process of heuristic generalization that Tsoukas, following Wittgenstein, proposes as the key to understanding. Finally, we reflect on how such an approach, despite (indeed precisely because of) its reluctance to center on the pursuit of generalizable, theoretical truths, could endow policymakers with that elusive ability that they so avidly crave: to make wise judgments about the design and implementation of *new* national eHealth programs. We conclude by discussing the significant challenges of engaging policymakers in the study of richness.

The Contested Place of the Case Study in Evaluating eHealth Programs

In two widely cited articles in Public Library of Science, a group of researchers, most of whom sat on the national steering group to evaluate the English NPfIT, offered a set of “methodologically robust” standards for evaluating eHealth programs (Catwell and Sheikh 2009; Lilford, Foster, and Pringle 2009). They proposed a quasi-experimental methodology in which the impact of eHealth programs is assessed as much as possible independently of their social and political context, for example, via systematic “step-wedge” designs in which later-adopting sites serve as controls for early adopters. Our own team took issue with these articles and proposed a diametrically opposing set of standards centered on in-depth case study (Greenhalgh and Russell 2010). We agreed with earlier scholars who depicted program evaluation not as experimentation but as social practice. We drew on Weick’s work on the “generative properties of richness” (thick description, reflexive theorizing, and “conceptual slack”—openness to the many new explanations that emerge when contextual detail is added to the account) in organizational case study (Weick 2007) and on the need to make collective sense when introducing technology in organizations (Weick 1990). We encouraged evaluators to immerse themselves in, and narratively describe, the social and political influences on eHealth programs.

Our alternative set of guiding principles met with a variety of responses. A swift riposte published in a leading health informatics journal, for example, exhorted researchers to “return to first principles” and reiterated the “known” benefits of the randomized controlled trial (Liu and Wyatt 2011). With academics so deeply divided, it is little wonder that policymakers were unsure what to make of “uncontrolled” case studies of eHealth programs, although given the extent to which they viewed the benefits of electronic records as self-evident (Markus and Keil 1994), it is possible that they saw no need for trials. More generally, however, while the experimental trial has a special place in the hearts of many doctors and while its epistemological foundations resonate strongly with the rationalist, “evidence-based” ideology that pervades contemporary policymaking (Greenhalgh and Russell 2009; Harrison 2002), few researchers with a background in social, political, or organizational science need to be persuaded of the merits, in principle, of in-depth case studies (sometimes referred to as *small-n* studies) when researching complex social programs.

Case Study: A Philosophical Taxonomy

To begin addressing the vexed question of how to move doctors and health care policymakers beyond a reductive “what works?” mind-set in eHealth programs, we must first take a detour into some ongoing philosophical debates *within* the community of case study researchers. The key question is: if we reject experimental and quasi-experimental studies on the grounds that they lack predictive power, how can we defend the case study against the countercharge of lack of generalizability (and the logical extension of this charge, that the richer a case study becomes, the less generalizable it will be)? How (if at all) can any level of prediction be achieved when the data set comprises a handful of descriptive studies, each of which is unique and unreplicable? To what extent (if at all) can the findings from such studies be used to inform program planning in *another* context or setting? How many cases are sufficient for a small-n sample? What is the significance (if any) of the sample of just one study?

Many authors before us have tried to summarize the large and contested literature on these questions. One of the clearest attempts is a book chapter proposing to classify different approaches to case study

research in terms of the differences in their underlying philosophical assumptions, not (as is more usually the case) in terms of the differences in methodology (Tsoukas 2009). Tsoukas acknowledges the positivist mainstream in his own discipline (organization and management studies). Here, at least until relatively recently, Newtonian models of reality dominated, and experimental (preferably, large-*n*) studies were viewed as the most robust route to generalizable truths. Tsoukas describes a post-Kuhnian shift from the pursuit of the “decontextualized ideal” to a recognition of the historical contingency of scientific claims, paralleled by a growing acceptance of (and, in some cases, a preference for) ethnographic and case study methods. But, he argues, while the *methodology* for studying complex social phenomena has moved on, many case study researchers (of whom probably the best known is Robert Yin) have retained an essentially experimental *epistemology* (Yin 1994).

Yin emphasizes the theoretical sampling of cases with the goal of analytic generalization (reasoning inductively through systematic cross-case comparison from a particular set of results to some broader theory of causation). Central to Yin’s methodology are (1) a selection of multiple cases, each of which is seen as representing a specific instance of the theoretical phenomenon being investigated; (2) the same types of data collected from each case in broadly the same way; (3) a detailed and methodical comparison of the cases’ specific features; and (4) rigorous testing of hypotheses concerning the relationships between the features. This approach to case study is preferred by many research sponsors and peer reviewers in the health care field, who tend to take their quality criteria from the experimental paradigm. But, Tsoukas argues, if taken to its logical conclusion, this approach would favor large-*n* samples, statistical testing of relationships between the variables, and articulation of the conclusions in terms of probabilistic reasoning.

Tsoukas suggests that at a philosophical level, case study research centers on the tension between two questions: “What is going on here?” (the study of the particular for its own sake) and “What is this a case of?” (the search for generalizability). Yin’s analytic generalization privileges the latter at the expense of the former, whereas case study researchers like Robert Stake, who favor naturalistic generalization (the learning that comes from the intrinsic study of the particular case) (Stake 1995), privilege the former, seemingly at the expense of the latter.

In health services research, there is currently much interest in realist evaluation, that is, the exploration via an in-depth case study of the

relationship of context, mechanism, and outcome (Pawson and Tilley 1997). This approach holds that the focus of research should be the “hard” (i.e., external and independent of our perceptions of it) social reality that forms the context for human action. Studying how interventions play out in this social reality, realists explain, helps elucidate the generative causality of social interventions: What [generally] works, for whom, in what circumstances? (Pawson and Tilley 1997). In realist notation, mechanisms are “underlying entities, processes, or [social] structures which operate in particular contexts to generate outcomes of interest” (Astbury and Leeuw 2010, 368). Although realist researchers often disagree on the fine detail of the context-mechanism-outcome relationship, they do share a preference for answering the general question, “What is this a case of?” to which the particular question, “What is going on here?” is subordinate. In an ideal world, Pawson and Tilley would recommend building and testing theory across a sample of cases. In practice, however, multiple comparative cases are hard to find, and Pawson subsequently developed realist review, a technique for looking retrospectively at single-case studies and generalizing by abstraction to produce theoretical insights (Pawson et al. 2005).

In sum, whereas the experimental researcher seeks to generalize via enumeration (*adding up* examples of the particular in large-*n* studies to produce *predictive* statements of general relevance), both Yin (analytic generalization from cross-case comparison) and Pawson and Tilley (realist evaluation) seek to generalize via theoretical abstraction (i.e., by *reasoning from* examples to produce *theoretical* statements of general relevance). This leaves unanswered the question on which the credibility of our own large, national-level evaluation hangs: What is the philosophical (and practical) significance of the “sample of one” study, especially when the researchers refrain from claiming that they are describing a case *of* anything? In other words, how can we justify an epistemology that privileges the particular question, “What is going on here?” over the general question, “What is this a case of?”

A Wittgensteinian View on the Singular Case Study

Ludwig Wittgenstein (1889–1951), whose life and work were elegantly summarized in Monk’s biography (1991), is perhaps best known as the

philosopher who changed his mind partway through a brilliant career. Indeed, Wittgenstein's change of mind was so dramatic that philosophy books invariably qualify descriptions of his work with "early" or "later." The young Wittgenstein, a protégé of Bertrand Russell, produced a formidable dissertation on the philosophy of logic, offering (among other things) a comprehensive theory of the logical structure of language. Soon afterward, in his home city of Vienna, Wittgenstein joined and strongly influenced the Vienna Circle of logical positivists. But after several years spent out of academia, he returned as a ruthless critic, not merely of logical positivism, but also of analytic philosophy more generally, and spared no invective for his own early work (but the rumor that he once attacked Karl Popper with a poker in a Cambridge debate is untrue).

The later (or "postanalytic") Wittgenstein became increasingly convinced that many things treated as philosophical problems and subjected to exhaustive analysis are inherently insoluble. In the words of his physician and philosophical disciple Maurice Drury, "However much the realm of what is explained is extended, the realm of the inexplicable is not reduced by one iota" (Drury 2000, 73). Wittgenstein's exhortation not to analyze all phenomena exhaustively inspired significant paradigm shifts, at least at the critical margins, in sociology (Winch 1958), anthropology (Geertz 2001), psychology (Harré and Tissaw 2005), and science and technology studies (Bloor 1983).

Rather than pursue the inherently fruitless holy grail of generalization by theoretical abstraction, Wittgenstein proposed that we should instead seek to understand the particular in all its unique, contextual detail. He was especially interested in the use of language. There *is* no formal system of the rules of language that accounts for every use of a word or phrase. Rather, words and phrases acquire particular meaning in particular situations, and language evolves as our use of it changes. For example, terms such as *leadership*, *multiprofessional team meeting*, and *consultation* mean different things in different countries, sectors, settings, departments, and even in the same department or team at different times. Wittgenstein used the term *language game* to depict the use of words (and also of gestures, behaviors, and so on) in social action (Wittgenstein 2009, para. 23). The rules of language, he declared, consist of no more than an open-ended and up-for-negotiation set of practices that are shared, to a greater or lesser extent, by individuals in a social group. It is through the study of the *active use* of language, and not in some abstracted book of grammar, that the "rules" become clear.

It follows that to understand a complex social situation, we must understand, first and foremost, the many and varied social practices of which it is made up, and the different (and perhaps conflicting) ways in which particular words, phrases, lists, instructions, taxonomies, gestures, behaviors, and the like are actually used by different groups. These practices are necessarily particular to the case. Indeed, it is only by grasping their contextual significance that we can understand the case. The extent to which we approach a case in the spirit of the later Wittgenstein is also the extent to which we must resist the temptation to begin with a closed definition of what “a case of X” comprises and then proceed to study how the case under investigation aligns with this (Tsoukas 2009).

It would be wrong to conclude from this summary that it is impossible to generalize from a Wittgensteinian case study analysis. Quite the contrary. However, the generalization is of a very specific kind. Wittgenstein’s notion of language games is not one that has no general rules but one whose rules are open-ended. An example is *consultation*. Physicians consult with patients; lawyers with clients; teachers with parents; planners with residents; and policymakers with citizens. Referring to only this last category, consulting may occur on soapboxes, in public meetings, through house-to-house surveys, on live-radio phone-ins, via Twitter, and so on—and within each of these formats are an infinite number of examples of how a “consultation” actually occurs, with differences in actors, contextual detail, framings, interactions, power dynamics, and so on. However diverse in form, consultations have things in common that are difficult to define in the abstract but are easy to recognize across a family of examples. These “family resemblances” can never be exhaustively defined because in each new case, the language game plays out in its own unique way (Wittgenstein 2009, para. 66–67).

If we were to research the social practice of “consulting with citizens,” exploring each case in order to understand its *particular* nuances, we would also enrich our *general* understanding of what it is to consult with citizens. More generally, case study research that focuses on what Wittgenstein called “the epistemology of the particular” works by expanding and sharpening the vocabulary and expressions as they are used by researchers, practitioners, policymakers, and citizens to talk about social practices, a process that Tsoukas calls *heuristic generalization* (Tsoukas 2009). In this way, they are able to draw ever more subtle distinctions between *this* instance of a particular social practice and *that* one.

The scholar who studies a painting of a tree by Cézanne and discusses with others its meaning and significance does not learn about merely this particular painting or the tree in it. Rather, he or she then will also look with a more sophisticated eye at other Impressionist paintings and other trees (Simons 1996). In a parallel process, case study evaluators collect and synthesize data and place their findings in the public domain in order to address the particular question, “What is going on here?” But through the reflection, discussion, and debate that (one hopes) ensues, they and their various audiences also achieve a greater understanding that can be applied more widely (Greenhalgh and Russell 2010).

A Case within a Case: The Newtown Integrated Records Pilot

As we described briefly here and in more detail elsewhere (Greenhalgh et al. 2010a, 2010c), the Summary Care Record and HealthSpace were (and, at the time of this writing, still are) large, complex eHealth programs that are part of the even larger and more complex, publicly funded NPfIT (Brennan 2007). To advance this article’s central argument—that immersion in detail is the route to understanding complex cases—we must use a small substudy to demonstrate how we arrive at heuristic generalization. Accordingly, we describe an attempt by a consultant diabetes specialist in Newtown (pseudonym), Dr. J., who was widely known as an innovator and enthusiast, to link his existing electronic patient record system with the HealthSpace system so as to allow his patients to view their own records, and the reciprocal attempt by Connecting for Health to use this unsolicited approach by a local champion to demonstrate the “interoperability” of its own state-driven record system with existing NHS legacy systems and private-sector add-ons that were running in parallel with it.

This case provided an interesting microcosm of the multiple stakeholders and institutional ambiguities associated with electronic patient records in the English NHS between 2007 and 2010. The Summary Care Record and HealthSpace were being introduced nationally, in staggered sequence, through the country’s 158 Primary Care Trusts (PCTs), whose numerous responsibilities included commissioning health care and supporting and overseeing local general practitioners (GPs).

Different GPs used different local record systems, which reflects the fact that historically, GPs were not employees of the NHS but independent contractors with it. Furthermore, despite the national mandate for NHS hospitals to use particular, centrally procured record systems, their development was years behind schedule, and the reality was that most hospitals were working with a variety of legacy systems that they had delayed upgrading in anticipation of the promised national solution (Robertson et al. 2010). In some areas, private suppliers and local entrepreneurs were moving in to fill the vacuum. Linknet (pseudonym) was a leading supplier of the NHS's electronic record systems, which in 2010 contained 18 million patient records, but it was not part of the NPfIT.

As in many towns in England at the time, in Newtown, parts of patients' hospital electronic records were accessible to their general practitioners before the NPfIT began. This record-sharing system was designed and introduced by Dr. J., whose long-standing interest in shared care, electronic records, and patient involvement (along with his personal charisma and considerable technical skill) had inspired a strong local vision for an electronic health record shared three ways: by hospital clinicians, primary care clinicians, and patients. The first step, which had not included patients, was to link the hospital system with general practitioners and community clinics (e.g., podiatry). All would share a core data set on patients with one of four long-term conditions (diabetes, coronary heart disease, chronic kidney disease, and stroke). The Newtown Integrated Record, as it was known, was launched in 2006 and became fully operational in 2009, when it began to be used regularly by the multidisciplinary foot clinic, prenatal clinic, diabetes clinic, and emergency teams at the hospital, and by primary and community care teams, each of which could call up a different customized view of the data held.

Patients' records held locally in participating organizations were linked using a "middleware" solution provided by Linknet. Middleware sits between applications that may be working on different operating systems, allowing one "incompatible" system to read from another. While in one sense Linknet's products competed with the systems introduced as part of the NPfIT, the company was cautiously positive about forming a partnership with Connecting for Health: "If we're going to have something like this [patient-accessible shared record], I'd prefer it if it was not owned by Microsoft, if it was neutral, so to speak, so in

principle I'm keen to work with HealthSpace" (senior executive, Linknet, SR22).

Dr. J. worked actively with the local Diabetes Patients' Forum to explore various possibilities for adding a patient portal to the system. The options included adapting an existing commercially available interface that had been designed for kidney patients, or a "paper solution," in which patients asked their GP to print out their record. Both options were rejected in favor of HealthSpace, partly because HealthSpace was considered to be a "more strategic" option, since it was the official NHS product and appeared to be the emerging market leader. Dr. J. approached Connecting for Health and was initially turned down, as this project was not part of its original business case, but Connecting for Health subsequently decided that Dr. J.'s proposal would be a worthwhile pilot that might open new avenues in other projects. Talks began in late 2008.

The Newtown Diabetes Patients' Forum strongly encouraged its members to sign up for the HealthSpace pilot, and eight (out of a clinic population of several thousand) registered for an advanced HealthSpace account in anticipation of joining the scheme. This process was complex, requiring the forum's members to take two separate proofs of identity to a front office in the PCT and then to decode a personal security access grid mailed to them. But when the patients tried to log on to their Newtown Integrated Record through HealthSpace, all but two of the eight got an error message:

After a couple of weeks I got the customary letter, and the card and what have you, and my password and all that was sorted out, I did all that, but like I say, the times I've tried to get onto it, it keeps coming up with the same thing, "Your GP isn't launched yet, your GP isn't taking part in this yet," is all it says. . . . and the surgery manager, she said, "Oh, I don't know nothing about that, I've never heard about it." (person with diabetes, SR04)

Despite the error message, this person's GP *was* participating in the pilot. The same researcher interviewed a PCT manager a few days later. He checked the system, found it "working," and suggested that the problem may be with the patient: "Well there's definitely data in the table for the Y— surgery, so I don't know why the person can't see that, maybe if they check their password, make sure they've done their permissions correctly" (PCT project manager, SR09).

Another patient with the same access problem contacted the HealthSpace help desk, which attributed the problem to either the GP or the local IT support at the PCT:

And they took it all on board and they were very nice, and then they said I'll ring you back, which they did, and they said this isn't a HealthSpace problem, this is, it looks like it's a local problem so it's probably, this is what they said, it's probably because either your GP hasn't uploaded his, your records onto the HealthSpace website yet, or it's something to do with the IT people on your local NHS area who are responsible for getting your doctor's records onto the HealthSpace site. (person with diabetes, SR11)

In the early stages of the pilot, two of the eight volunteer patients managed to log onto the system and access their Newtown Integrated Record via HealthSpace. They valued the information and liked the design of the record, which they found easy to navigate and interpret. For a short time, the "tipping point" seemed imminent.

The project continued on a very small scale for several months, even though only two patients ever accessed their Newtown Integrated Records (this near-static phase was attributed to changes in staff at the PCT). A server upgrade was planned for mid-2009 to increase capacity before widening the scheme. Then, unexpectedly, a "bug" appeared in the new server that made the system run slowly from the clinicians' end and blocked the patients' access entirely. Tests by a subcontractor to Connecting for Health attributed the bug to the upgrade's "browser incompatibility," so Connecting for Health decided to turn off the link between HealthSpace and the Newtown Integrated Record until this had been fixed.

We worked with Linknet and Newtown PCT to try and resolve what the issue was. We did spend a few weeks trying to get to the bottom of where the issue lay. And then, until mid-July [2009], and Linknet, they basically acknowledged that further testing wasn't going to achieve anything and that they would need to implement a fix within their system. . . . So the work that has been going on since mid-July has been, well basically, waiting for Linknet to implement this fix. And also waiting for additional test accounts to be created so we could test the performance level of the records. (staff member, HealthSpace team, Connecting for Health, SR17)

From the clinicians' perspective, the key problem was how the Linknet middleware was interfacing with the HealthSpace application. The diabetes center had contracted with Linknet to create dummy patients on the new Newtown Integrated Record system and undertake security testing, but that work could not proceed until the HealthSpace link was turned on again. "Basically it seems to be a Linknet/Connecting for Health axis that is required to resolve it" (hospital consultant, SR20).

Connecting for Health's arrangement with Linknet appeared to preclude direct dialogue: "The relationship is between the PCT and Linknet, we don't have a relationship with Linknet as such, in that we don't have any contractual relationship with them. The contractual relationship would be between the PCT and Linknet" (staff member, HealthSpace team, Connecting for Health, SR17). From Linknet's perspective, the main problem was the stringent standards and testing requirements imposed by Connecting for Health and the associated costs, which had not been anticipated by all parties when the project was set up. Connecting for Health required "penetration testing" (i.e., rigorous security testing to see whether the system could be hacked into) but considered that the cost of this should be borne by the PCT and/or Linknet. "As far as we're concerned, we've done everything we need to do, and it's back to them. Is it overregulation? Is it overtesting? Probably a bit of both. We've done our bit ages ago, and for some reason it's not moving ahead" (senior executive, Linknet, SR22).

Different stakeholders in this small but complex subproject had different views of what exactly was being "tested." In reality, it was not the individual, static components of the system that needed testing but how the live system functioned *dynamically* in real time. But tellingly, each stakeholder tended to perceive his or her own component to be "working fine" and expressed suspicion about those components for which other parties were responsible:

The pen[etration] testing people came along, and classically communications errors got in the way. The people from Z— [a small private company, subcontracted to Connecting for Health] who came to do the testing weren't briefed properly, and we weren't either. We said to them, we need to understand exactly what you want to do and if you think we should do something more around security. We were expecting them to access HealthSpace and make sure it was secure, see if it was all working. But they assumed they were coming to look at *our* end, not check *their* end. . . . It seemed reasonable, but they then

wanted to test the live system, it's in operation across the patch and we didn't want them jeopardizing the whole system as it's up and running. I was happy to let them have access to HealthSpace but then they started up with "We just want to let this tool run on your network," and we said, "Yes, well, no, thank you, it's a live network," and then we thought we didn't realize what they had been told to do was to be targeted towards a live system, and they hadn't realized either. So they went away again. (senior manager, Newtown PCT, SR23)

The HealthSpace–Newtown Integrated Record project had been created by a local clinician and his patients' enthusiasm. It had a somewhat "ad hoc" status; for example, it was not using the project management tools and processes that were in standard use for "official" technology projects within the PCT. Perhaps for this reason, an IT manager at the PCT, who had been identified by HealthSpace staff as the central contact for the project locally, saw this work as a "side issue" and the staff's own involvement in it as peripheral.

Because the HealthSpace part, from our perspective, was never a formal project; if it had been a formal project, if it had been part of Newtown Integrated Records, PID [project initiation document] and business case, that we were formally integrating HealthSpace then. . . . It's always been kind of like a side thing. I'm only vaguely involved—I just see the emails. . . . so [name] like just kept me up to date with what was going on really. What's happened is you've got somebody who's very enthusiastic and wants to implement it but then not gone down formal channels and obviously there's a clinician who's very keen and you can see the benefits of doing it for his patients but then fundamentally it then gets undermined. (PCT IT manager, SR19)

The preceding comment is particularly ironic given the overall emphasis within the NPfIT on "benefits realization" and "engaging local champions" and the early indications that patients greatly valued access to their Newtown Integrated Record. In our 2010 report, we concluded that "despite enthusiasm from patients, clinicians, the PCT, Linknet and Connecting for Health, and a great deal of work that was undertaken by all parties, challenges relating to information governance and complex commercial relationships are yet to be overcome in this promising project" (Greenhalgh et al. 2010b, 179). A year later, little has changed and patients' access to their records in Newtown remains, in Latour's words, a technological fiction (Latour 1996).

Analysis

This account of one clinical team's efforts to link an existing electronic record system to the NPfIT is lengthy but, at the same time, compressed. Much detail is given, but because complex cases sit within wider cases and beg ever richer contextual explanations and because ambiguity demands to be teased out and explored, much is also excluded. Our account also has a classic narrative form, with the local clinical champion depicted as struggling heroically to improve care and empower his patients against an inflexible, bureaucratic, and technically imperfect system. Dr. J.'s efforts appeared to be stymied by a complex combination of circumstances that we were never able to explain. Even though there were no obvious villains, there were bystanders whose lack of engagement in this project (for whatever reason) did have a bearing on the story. All this human action and inaction took place against a backdrop of competing priorities and conflicting value systems among participating stakeholder organizations, and in the context of software products that "worked" individually but not in combination. In addition, the social and technical dynamics were opaque not only to the readers of the case study but also to its central characters.

Some of the usual tensions in eHealth implementation (e.g., among local legacy systems, private-sector "solutions," and a nationally procured standard product; or between a local innovator and an inflexible central protocol) are evident, so we could offer this story as a "case of" such generic issues and find other cases in order to take such analytic generalization further. But this approach would miss the point that the narrative is striking mainly for its unique and puzzling features. The key finding is not that the individual actors or organizations in the story are illustrative of *generic* problems or issues but that *in this particular instance* they never achieve a shared perspective on the project and hence do not manage to pull together sufficiently to make it happen. Focusing on the detail of the fragment thus allows us to shift the crucial Wittgensteinian frame from "a dead, mechanically connected world" (i.e., a world of abstracted variables such as "leadership," "innovation," and "resistance") to "a living world of responsive relations" (i.e., a story about a particular leader, a particular innovation, and particular people who resisted its introduction) (Shotter and Tsoukas 2011, 320). This shift, some have argued, is crucial when studying the implementation of information systems in organizations, because of the limitations of "variance models" to

capture their uniqueness and sociotechnical interdependencies (Ramiller and Pentland 2009).

The rich narrative form and detailed ethnographic basis of this case fragment show how actors from different social, organizational, and cultural worlds (physician, technician, public-sector manager, private-sector executive, subcontractor, patient) pass by one another instead of reaching what interpretivist information systems researchers would call accommodation (Checkland and Holwell 1998). Many actors in this story, for example, appear to hold static, oversimplified models of technical systems (they assume that components can be tested individually and passed as “working,” and they do not anticipate that an upgrade to one component may have a *downgrading* impact on the system as a whole). This leads them to imagine a system in which “our” components work fine but “their” components do not function (or are not operated correctly) and to frame the solution in terms of a [technical] “fix.”

The actors in this case fragment bring different professional and institutional perspectives to questions of security and information governance, which are brought to bear dynamically, in the here-and-now, as the action unfolds. It is not that anyone disagrees *in the abstract* about what the security standards are. Rather, the question is whether *this* subcontractor may be permitted access to *this* system, having turned up *today* with an ambiguous brief.

To the extent that this case study succeeds, it does so mainly by highlighting the numerous language games that were being played in Newtown (contracting, collaborating, negotiating, consulting, security testing, piloting, legitimating, justifying, and so on), each role taken by actors from different social worlds who drew on (at best) partial understandings of one another’s rules and meaning systems. This in turn led to recurring situations in which both individuals and the organizations they represented found themselves talking and acting at cross-purposes. These real-time misunderstandings perhaps were exacerbated by the number of documents in internal circulation that were marked “commercial in confidence” and that were pointedly *not* shared among the different social worlds.

The failure to achieve collective sensemaking that characterized this case fragment was also evident across the NPIIT more generally. Not only junior staff but also chief executives, lead clinicians, and senior designers—indeed, at least 130 of the 140 individuals we interviewed—admitted openly to not fully understanding what was going on. A

handful of individuals (some employed by Connecting for Health; others in the IT industry; and one or two clinician champions) appeared to grasp the programs' numerous complexities. These people were technically skilled, politically and commercially savvy, and (usually) clinically trained. They moved flexibly among the different worlds of policy, clinical practice, health care management, the commercial IT industry, and the patient and public spheres. Across these worlds, they built a multifaceted grasp of the programs and worked to help others make sense of them.

But nobody understands everything or has the time to find out about, and keep pace with, all the components and interactions in a complex system. Those aspects of the Summary Care Record and HealthSpace programs that failed could sometimes be traced back to a mission-critical loss of overview (or, sometimes, a change of job) by one of these pivotal individuals. Even if such a person had been found to help make sense of the Newtown subproject (Dr. J. came close, but he was relatively naïve about the NPfIT's politics), the program as a whole was so vast and complex that crucial misunderstandings somewhere in the system were inevitable.

Discussion

As Weick observed, "Richness restrains hubris" (Weick 2007, 18). The complexity of contemporary health care, combined with the multiple stakeholders and perspectives in large technology initiatives, means that national eHealth programs require considerably more thinking through than has sometimes been the case to date. This article argued that the rich descriptions made possible by in-depth case study are the key to understanding the dynamic complexities of such programs. Using insights from Wittgenstein's postanalytic philosophy and an illustrative case fragment, we have suggested that such programs can be conceptualized as a series of overlapping, conflicting, and mutually misunderstood language games that typically combine to produce a situation of ambiguity, paradox, incompleteness, and confusion for which technical fixes offer only the appearance of a solution. This reality jars uncomfortably when the policy vision consists of a "state-of-the-art" system that will be fully and unproblematically implemented according to the original blueprint, if only all people do what they are contracted and expected to do.

Critical academics have proposed that the introduction, implementation, and evaluation of eHealth programs inevitably reflect and perpetuate the wider alignments of political power, which Michel Foucault called “régimes of truth” (Introna 2003). But it is arguably not necessary to view the world through a Foucauldian lens to accept that when policymaking takes a rationalist turn—shifting from deliberative to rule-based decision making, focusing narrowly on the pursuit of “what works,” valuing managerialism over professionalism, and introducing an ever tighter surveillance of performance—it becomes almost impossible to articulate a national eHealth program as anything other than a detailed advanced specification with firm milestones and carefully delineated work packages. In such contexts, those working to implement policy have little choice but to view the key task as controlling, coordinating, and aligning these various packages rather than, for example, understanding and accommodating the various nuanced language games being played by different stakeholders.

It follows from these (flawed) rationalist assumptions that a national eHealth program *is* the sum total of the predefined work packages occurring across the country and that knowledge of “what is going on here” can be equated with a central “dashboard” of progress metrics aggregated from the weekly reports submitted by regions and localities. When commentaries on the program evince confusion, the task for senior policymakers would appear to be to swiftly resolve the problems in order to place the program back on track. Within days of the publication of our report, *The Devil's in the Detail* (Greenhalgh et al. 2010b), precisely this chain of events had been set in motion, generating a set of in-house reports that, in the short term at least, simplified and contained the redefined problem. Note that from the policymakers' perspective, the in-depth case study not only has low marginal utility but perhaps also a *disutility*, in that acquiring detailed and nuanced knowledge of a complex program could make the various language games (and, indeed, other political games) more difficult to play (Scharpf 1991). Because national eHealth programs stretch the limits of comprehension, it perhaps is understandable for policymakers to seek to simplify their world and exhibit “satisficing behavior” within it, that is, to ensure that their decisions make sense and are accountable within selected parameters (Simon 1965). With rationality thus bounded (Gigerenzer 2002), the quasi-experimental study of “what works” is an alluring research design.

The findings of this study, however, suggest that if the dismal track record of national eHealth programs is to improve, some policymakers (and academics who share their desire to rationalize, simplify, and create a unified order) must be helped to escape from this frame of reference. They need to be encouraged to contemplate, for example, that at a high-level stage of design, a modular approach based on a few core principles (e.g., a common user interface, local interoperability) may be more useful and enduring and, paradoxically, a *less* risky investment option than an overarching blueprint, tight specification, and inflexible master plan. Indeed, an optimist might extend this argument to conclude that huge waste could be avoided by using philosophical insights to justify excluding from investment options these inflexible and monolithic eHealth programs, however “state of the art” they may be.

Policymakers also need help to escape from the assumption that at the operational level, they can make collective sense of “equivocal” technologies (Weick 1990) through briefing breakfasts that present a single, approved, official version of events and priorities. Instead, “clinical engagement,” as well as “manager engagement,” “patient engagement,” and “industry engagement,” may be more productively viewed as a necessarily conflict-ridden process requiring ongoing resources and efforts. As we concluded in our final report on the Summary Care Record and HealthSpace (Greenhalgh et al. 2010b, 21):

Dialogue (or lack of it) occurs in the context of multiple conflicting worlds (political, clinical, technical, commercial, academic and personal—and probably others as well). Strong feelings, misunderstandings, conflicting values and competing priorities are to be expected—and we offer no magic recipe for resolving them. But we do offer an observation from three years’ involvement with these complex programmes: greatest progress appeared to be made when key stakeholders came together in uneasy dialogue, speaking each other’s languages imperfectly and trying to understand where others were coming from, even when the hoped-for consensus never materialised.

Our findings also suggest that the individual change agent could gain much by employing people with long and diverse experience as “cultural brokers” between these different worlds. The ideal leader to help us “muddle through” a national eHealth program’s ambiguities and paradoxes would appear to be a clinician who is both technologically and politically astute, has spent time in academia and industry, still works at

least part-time in clinical practice, has a mandate from a technology user group to represent their perspectives, and is also comfortable working in a committee environment. Given how much appears to hang on the input of such individuals, resources should be invested in their identification, recruitment, retention, and support in a way that goes beyond a reductive focus on their “competencies.”

Let us revisit the four possible explanations we introduced earlier for why our final evaluation report appeared to be unpopular with policy-makers. They may have judged our work to be of poor quality because we did not attempt to resolve the program’s many ambiguities and paradoxes (indeed, we meticulously exposed them and declared them insoluble). The Summary Care Record and HealthSpace may have become runaway technology projects because the policymakers’ reluctance to engage with the programs’ messiness and situated detail allowed them to continue to anticipate that a “tipping point” was imminent. The programs’ high degree of path dependency was partly inherent in their vast scale, but it was made worse by an adherence to prespecified procedures and milestones that left little room for debate or deliberation. All these eventualities might have been avoided if the senior civil servants who set the direction and pace of the programs and briefed the ministers had had the rich vocabulary and nuanced understandings that would have allowed them to engage in what has been called “frame-reflective policy practice” (Rein and Schön 1996).

A reviewer of an earlier draft of this article pointed out that our analysis has a “Russian doll” aspect. That is, the Newtown case fragment serves as a window to the HealthSpace program, which in turn highlights many of the wider problems of the NPfIT in England. The heuristic generalization made possible by teasing out what was going on in Newtown also helps reveal the widespread nonadoption and abandonment of another much larger component of the NPfIT, the Detailed Care Record Scheme, which sought to replace paper records with full electronic records in general practice and in hospitals, at an estimated cost of £7 billion (US\$11 billion) (National Audit Office 2011). More broadly still, it illustrates a number of challenges that are likely to play out to a greater or lesser extent in all large-scale eHealth programs. Its lessons strike a note of caution, for example, regarding efforts to introduce telemedicine and remote monitoring on a regional or national scale, which might similarly be considered as a set of overlapping language games among the clinical, technical, political, commercial, and personal worlds.

Our “sample of one” case study differs in many respects from the other national electronic record programs to which we referred briefly in our introduction. As the political analyst Carolyn Tuohy observed, the fortunes of Britain’s health care system have tended to turn on a “big bang” logic in which a powerful state periodically introduces large-scale reforms relatively quickly and policy tends to be driven by big plans (and big “accidents”) rather than by a policy process more attuned to local contextualities and learning from mistakes (Tuohy 1999). This was the context in which the NPfIT emerged as a massive exercise in project management with nonnegotiable goals and a short timescale, requiring an almost unprecedented degree of satisficing on the part of policymakers.

The U.S. health care system, in contrast, has a “mosaic” logic, in which reform tends to occur more patchily (on a regional rather than national basis) and with fewer dramatic turns (Marmor and Oberlander 2011; Tuohy 1999). The emphasis of the HITECH scheme on outcomes (i.e., encouraging the achievement of “meaningful use” metrics), rather than on the nationwide procurement and implementation of a particular information system, may suggest a greater use of local and/or national learning communities that might achieve an understanding of overlapping language games (Friedman, Wong, and Blumenthal 2010). It does not, however, mean that this emergence is inevitable. Critics of HITECH have argued that “meaningful use” is a restrictive, top-down policy that rests on deterministic assumptions, flawed logic, and reductive metrics of success (Hussain 2011). HITECH may favor carrots over sticks and outcomes over processes, but it is arguably a far cry from the explicit emphasis on collaborative learning, “middle-out” development model and continually renegotiated goals characteristic of New Zealand’s highly successful national electronic record system (Bowden 2011; Greenhalgh and Bowden 2010). At the time of this writing, the only thing that we can confidently say about HITECH is that a detailed empirical case study is likely to generate many more lessons for the international community and enrich our understanding of eHealth policymaking more generally.

While it is possible to theorize national eHealth programs with the goal of analytic generalization (indeed, we have offered such theorizations ourselves; see Greenhalgh and Stones 2010), this may be less practically useful at the policy level than the approach described here. That approach is to achieve through the detailed study of particular programs the ability

to reflect and deliberate on the complex, interdependent social practices that make such programs unique. A scholarly analysis of the case-of-one is engaged and interpretive rather than abstracted and representational. It depends on what has been called “theory’s ineradicable dependence on the dynamics of the life-world within which it has its ‘currency’” (Shotter and Tsoukas 2011, 311).

We raise the question, “Of all dead (and living) philosophers, why Wittgenstein?” While a perspective based on the work of this particular philosopher has allowed us to reach the preceding conclusions, we do not claim that this is the only lens through which this contested terrain might be studied. Reviewers and colleagues have suggested that similar, though not identical, conclusions might be reached by building on the work of scholars such as Gottlob Frege, Richard Rorty, Irving Goffman, Harold Garfinkel, and Charles S. Peirce.

Whatever tribal colors they choose, evaluators of eHealth programs face a number of challenges. First, they need to engage with all aspects of such programs in order to produce richly illuminative case studies that will open up and inform the debate about what is going on. Second, they need to develop ways of drawing judiciously on these case studies and working *within* policymakers’ bounded rationality, perhaps using scenarios or other heuristic devices, to inform and influence what the latter have chosen, perhaps for good reason, not to understand. Third, they need to transcend, and help others transcend, the linear and deterministic metaphors (“research into practice”) implying that the core task is determining what works, after which implementing an eHealth technology will be a straightforward exercise in project management.

Large-scale eHealth programs have an almost universally negative track record, with the exception of a handful of single-organization schemes such as Kaiser Permanente and the U.S. Veterans Health Administration (Shekelle and Golzweig 2009). If this record is to improve, all stakeholders must recognize that they each bring only one frame of reference and only one set of language games to the learning community. They must acknowledge the existence and the legitimacy in context of language games following the different rules of other social worlds. And they must commit to the messy deliberative process that, precisely because its cacophony of perspectives will be revealed as incommensurable, offers the best journey toward accommodation and moving forward. That, as Jack Nicholson might have said, will be as good as it gets.

References

- Astbury, B., and F. Leeuw. 2010. Unpacking Black Boxes: Mechanisms and Theory Building in Evaluation. *American Journal of Evaluation* 31:363–81.
- Bloor, D. 1983. *Wittgenstein: A Social Theory of Knowledge*. New York: Columbia University Press.
- Blumenthal, D., and M. Tavenner. 2010. The “Meaningful Use” Regulation for Electronic Health Records. *New England Journal of Medicine* 363(6):501–4.
- Bowden, T.C. 2011. EHR Strategy: Top Down, Bottom Up or Middle Out? *Studies in Health Technology and Informatics* 164:138–42.
- Brennan, S. 2007. The Biggest Computer Programme in the World Ever! How’s It Going? *Journal of Information Technology* 22:202–11.
- Brown, N. 2007. Shifting Tenses—From Regimes of Truth to Regimes of Hope? *Configurations* 13(3):331–55.
- Catwell, L., and A. Sheikh. 2009. Evaluating eHealth Interventions: The Need for Continuous Systemic Evaluation. *PLoS Med* 6(8):e1000126.
- Cayton, H. 2006. Report of the Ministerial Taskforce on the NHS Summary Care Record. Available at www.connectingforhealth.nhs.uk/resources/care_record_taskforce_doc.pdf (accessed August 8, 2011). London: Connecting for Health.
- Checkland, P., and S. Holwell. 1998. *Information, Systems, and Information Systems: Making Sense of the Field*. Chichester: Wiley.
- Collins, T. 2007. Government Snubs Key Recommendations of Public Accounts Committee on NHS National Programme for IT. *Computer Weekly*. Available at <http://www.computerweekly.com/blogs/public-sector/2007/08/government-snubs-key-recommend.html> (accessed July 30, 2011).
- Collins, T. 2010. NPfIT—Business as Usual? The Death of the NPfIT Has Been Exaggerated. *Computer World UK*. Available at <http://blogs.computerworlduk.com/the-tony-collins-blog/2010/09/npfit—business-as-usual/> (accessed July 30, 2011).
- Department of Health. 2005. *Connecting for Health: A Guide to the National Programme for Information Technology*. London: HMSO.
- Department of Health. 2010. *Liberating the NHS: An Information Revolution*. London: HMSO.
- Drury, M.O. 2000. *The Danger of Words and Writings on Wittgenstein*. Ed. D. Berman, M. Fitzgerald, and J. Hayes. London: Thoemmes Press.
- Fisher, E. 2010. Contemporary Technology Discourse and the Legitimation of Capitalism. *European Journal of Social Theory* 13(2):229–52.

- Friedman, C.P., A.K. Wong, and D. Blumenthal. 2010. Achieving a Nationwide Learning Health System. *Science Translational Medicine* 2(57):57cm29.
- Geertz, C. 2001. *Available Light: Anthropological Reflections on Philosophical Topics*. 2nd ed. Princeton, NJ: Princeton University Press.
- Gigerenzer, G. 2002. *Bounded Rationality: The Adaptive Toolbox*. Cambridge, MA: MIT Press.
- Greenhalgh, T., and T. Bowden. 2010. Why Many Hands Make IT Work. *Health Service Journal* 120(6224):12–13.
- Greenhalgh, T., S. Hinder, K. Stramer, T. Bratan, and J. Russell. 2010a. Adoption, Non-Adoption and Abandonment of an Internet-Accessible Personal Health Organiser: Case Study of HealthSpace. *BMJ* 341:c5814.
- Greenhalgh, T., and J. Russell. 2009. Evidence-Based Policymaking: A Critique. *Perspectives in Biology and Medicine* 52(2):304–18.
- Greenhalgh, T., and J. Russell. 2010. Why Do Evaluations of eHealth Programs Fail? An Alternative Set of Guiding Principles. *PLoS Med* 7(11):e1000360.
- Greenhalgh, T., and R. Stones. 2010. Theorising Big IT Programmes in Healthcare: Strong Structuration Theory Meets Actor-Network Theory. *Social Science & Medicine* 70:1285–94.
- Greenhalgh, T., K. Stramer, T. Bratan, E. Byrne, J. Russell, S. Hinder, and H.W.W. Potts. 2010b. *The Devil's in the Detail: Final Report of the Independent Evaluation of the Summary Care Record and HealthSpace Programmes*. London: University College.
- Greenhalgh, T., K. Stramer, T. Bratan, E. Byrne, J. Russell, and H.W.W. Potts. 2010c. Adoption and Non-Adoption of a Shared Electronic Summary Record in England: A Mixed-Method Case Study. *BMJ* 340:c3111.
- Harré, R., and M. Tissaw. 2005. *Wittgenstein and Psychology: A Practical Guide*. Farnham: Ashgate.
- Harrison, S. 2002. New Labour, Modernisation and the Medical Labour Process. *Journal of Social Policy* 31(465):485.
- Hussain, A.A. 2011. Meaningful Use of Information Technology: A Local Perspective. *Annals of Internal Medicine* 154(10):690–92.
- Introna, L.D. 2003. Disciplining Information Systems: Truth and Its Regimes. *European Journal of Information Systems* 12:235–40.
- Kreps, D., and H. Richardson. 2007. IS Success and Failure—The Problem of Scale. *Political Quarterly* 78(3):439–46.
- Latour, B. 1996. *Aramis, or Love of Technology*. Trans. Catherine Porter. Cambridge, MA: Harvard University Press.

- Lilford, R.J., J. Foster, and M. Pringle. 2009. Evaluating eHealth: How to Make Evaluation More Methodologically Robust. *PLoS Med* 6(11):e1000186.
- Liu, J.L., and J.C. Wyatt. 2011. The Case for Randomized Controlled Trials to Assess the Impact of Clinical Information Systems. *Journal of the American Medical Informatics Association* 18(2):173–80.
- Markus, L.M., and M. Keil. 1994. If We Build It, They Will Come: Designing Information Systems That People Want to Use. *Sloan Management Review* 35(4):11–25.
- Marmor, T., and J. Oberlander. 2011. The Patchwork: Health Reform, American Style. *Social Science & Medicine* 72(2):125–28.
- Monk, R. 1991. *Wittgenstein: The Duty of Genius*. London: Penguin.
- National Audit Office. 2011. *The National Programme for IT in the NHS: An Update on the Delivery of Detailed Record Systems*. London: HMSO.
- Pawson, R., T. Greenhalgh, G. Harvey, and K. Walshe. 2005. Realist Review—A New Method of Systematic Review Designed for Complex Policy Interventions. *Journal of Health Services Research & Policy* 10(suppl. 1):21–34.
- Pawson, R., and N. Tilley. 1997. *Realistic Evaluation*. London: Sage.
- Ramiller, N., and B. Pentland. 2009. Management Implications in Information Systems Research: The Untold Story. *Journal of the Association for Information Systems* 10(6):474–94.
- Rein, M., and D. Schön. 1996. Frame-Critical Policy Analysis and Frame-Reflective Policy Practice. *Knowledge, Technology & Policy* 9(1):85–104.
- Robertson, A., K. Cresswell, A. Takian, D. Petrakaki, S. Crowe, T. Cornford, N. Barber, A. Avery, B. Fernando, A. Jacklin, R. Prescott, E. Klecun, J. Paton, V. Lichtner, C. Quinn, M. Ali, Z. Morrison, Y. Jani, J. Waring, K. Marsden, and A. Sheikh. 2010. Implementation and Adoption of Nationwide Electronic Health Records in Secondary Care in England: Qualitative Analysis of Interim Results from a Prospective National Evaluation. *BMJ* 341:c4564.
- Sarup, D. 2003. “To Be, or Not to Be”: The Question of Runaway Projects. *Information Systems Control Journal* 5. Available at http://www.softwareresearch.net/fileadmin/src/docs/teaching/SS06/PM/runawayprojects_saca.pdf (accessed August 9, 2011).
- Scharpf, F.W. 1991. Games Real Actors Could Play: The Challenge of Complexity. *Journal of Theoretical Politics* 3:277–304.
- Shekelle, P.G., and C.L. Goldzweig. 2009. *Costs and Benefits of Health Information Technology: An Updated Systematic Review*. London: Health Foundation on behalf of the Southern California Evidence-Based Practice Center, RAND Corporation.

- Shotter, J., and H. Tsoukas. 2011. Theory as Therapy: Wittgensteinian Reminders for Reflective Theorizing in Organization and Management Theory. *Research in the Sociology of Organizations* 31:311–42.
- Simon, H.A. 1965. *Administrative Behaviour: A Study of Decision-Making Processes in Administration Organisation*. New York: Free Press.
- Simons, H. 1996. The Paradox of Case Study. *Cambridge Journal of Education* 26(2):225–40.
- Stake, R. 1995. *The Art of Case Study Research*. London: Sage.
- Tsoukas, H. 2009. Craving for Generality and Small-n Studies: A Wittgensteinian Approach towards the Epistemology of the Particular in Organization and Management Studies. In *The Sage Handbook of Organizational Research Methods*, ed. D.A. Buchanan and A. Bryman, 285–301. London: Sage.
- Tuohy, C.J. 1999. *Accidental Logics: The Dynamics of Change in the Health Care Arena in the United States, Britain, and Canada*. New York: Oxford University Press.
- Weick, K.E. 1990. Technology as Equivoque: Sensemaking in New Technologies. In *Technology and Organizations*, ed. P.S. Goodman and L.S. Sproull, 1–44. San Francisco: Jossey-Bass.
- Weick, K.E. 2007. The Generative Properties of Richness. *Academy of Management Journal* 50(1):14–19. Available at http://www.people.vcu.edu/~rsleeth/MGMT691/Supplement/Weick_on_Dropping_Tools_and_Ethnographic_Richness.pdf (accessed August 9, 2011).
- Westbrook, J.I., and J. Braithwaite. 2010. Will Information and Communication Technology Disrupt the Health System and Deliver on Its Promise? *Medical Journal of Australia* 193(7):399–400.
- Whitt, M.R. 2010. Federal Auditor General's Report on Canada Health Infoway. *Health Law in Canada* 30(3):145–46.
- Wilsford, D. 1994. Path Dependency, or Why History Makes It Difficult but Not Impossible to Reform Health Care Systems in a Big Way. *Journal of Public Policy* 14:251–83.
- Winch, P. 1958. *The Idea of a Social Science and Its Relation to Philosophy*. London: Routledge & Kegan Paul.
- Wittgenstein, L. 2009. *Philosophical Investigations*. 4th ed. Trans. P.M.S. Hacker and J. Schulte. Oxford: Wiley-Blackwell.
- Yin, R.K. 1994. *Case Study Research: Design and Methods*. 2nd ed. London: Sage.

Acknowledgments: This article draws heavily on work by Haridimos Tsoukas (see references), who, to our knowledge, was the first to apply Wittgenstein's ideas to organizational case study. Our own, more limited, original contribution was to apply Tsoukas's reading of Wittgenstein to the study of policymaking. We also acknowledge both the fieldwork by other members of the original evaluation team for the Summary Care Record and HealthSpace programs (Katja Stramer,

Tanja Bratan, Susan Hinder, Henry Potts, and Emma Byrne), which contributed to the data set analyzed here, and our conversations with numerous colleagues as the ideas expressed here took shape. Detailed and constructively critical comments from three reviewers (Martin Connor, Scott Silverstein, and John Gabbay) and the editor (Brad Gray) enabled us to improve significantly the initial draft of this article.