Available online at www.sciencedirect.com

ScienceDirect

Research Policy 37 (2008) 330–352

research policy

www.elsevier.com/locate/respol

Galileo's stream: A framework for understanding knowledge production

J.Y. Tsao*, K.W. Boyack, M.E. Coltrin,
J.G. Turnley, W.B. Gauster

Sandia National Laboratories, P.O. Box 5800, Albuquerque, NM 87185-0601, United States

Received 15 December 2004; received in revised form 21 December 2006; accepted 2 October 2007
Available online 20 December 2007

Abstract

We introduce a framework for understanding knowledge production in which: knowledge is produced in stages (along a research to development continuum) and in three discrete categories (science and understanding, tools and technology, and societal use and behavior); and knowledge in the various stages and categories is produced both non-interactively and interactively. The framework attempts to balance: our experiences as working scientists and technologists, our best current understanding of the social processes of knowledge production, and the possibility of mathematical analyses. It offers a potential approach both to improving our basic understanding, and to developing tools for enterprise management, of the knowledge-production process.
Published by Elsevier B.V.

Keywords: Knowledge production; S&T policy; R&D; Paradigm creation; Disruptive innovation

1. Introduction

This work is an outgrowth of a strategic-planning exercise in which we asked the question: how can a small science-centric organization best contribute to a larger technology-centric research-and-development laboratory of which it is part. As we initiated this exercise, we surveyed existing frameworks within which we might debate and answer that question. None were found satisfactory. We did find, however, potential *elements* for such a framework from two communities asking related questions.

A first potential element stemmed from the ideas of Kuhn (1962) and other philosophers of science, who viewed knowledge production as proceeding in stages: "revolutionary" science creating new paradigms for understanding the world, and "normal" science extending those paradigms and imbuing them with richness and depth. These stages conveyed a sense of trajectory, and one could easily imagine that, in emphasizing different portions of such trajectories, a science-centric organization would contribute differently to the technology-centric laboratory of which it is part.

A second potential element stemmed from the ideas of Bush (1945), Stokes (1997) and other knowledge production policy-makers or political scientists, who viewed science as embedded in a larger national (and even global) science and technology environment. Science could be produced with various degrees of interactivity with that environment, and one could easily imagine that, in emphasizing different degrees of interactivity, a science-centric organization would again contribute

---

* Corresponding author. Tel.: +1 505 844 7092;
fax: +1 505 844 4045.
E-mail address: jytsao@sandia.gov (J.Y. Tsao).

0048-7333/$ – see front matter. Published by Elsevier B.V.
doi:10.1016/j.respol.2007.10.004

differently to the technology-centric laboratory of which it is part.

Both of these elements were reasonable, and accepted to some degree within their own communities. They had not yet, however, been combined as elements of a larger and more practically useful framework for knowledge production. Hence, we took a detour (rather significant, as it turns out) to develop a framework based on just such a combination and on following that combination to some of its logical conclusions.

We call this framework "Galileo's stream." That name is intended to convey the sense in which we are generalizing the work of Donald Stokes, who coined the phrase "Pasteur's quadrant" (Stokes, 1997), to include more varieties of knowledge production. It is also intended to convey a sense of the importance of these other varieties, one of which we call Galileo's stream, in which new technology enables breakthroughs in science.

Although one might view this new framework as simply a generalization of Pasteur's quadrant, one might also view this generalization as *necessary* for a more fine-grained, or "micro," description of knowledge production. Indeed, we are cautiously optimistic that this more fine-grained description might someday enable a quantitative dynamical theory of knowledge production. Such a theory could supplement and rationalize a large and growing body of observations and case histories of knowledge production processes (see, e.g., Freeman and Soete, 1997 and Mokyr, 2002). And such a theory could perhaps eventually be linked with economic models of growth in which knowledge and ideas play a central role (Romer, 1994; Jones, 2002; Tang, 2005).

The Galileo's stream framework is organized into three "layers," which we discuss in turn.

In Section 2, we introduce the foundational layer of the framework. We call this layer the "knowledge-production space" layer because its central feature is that knowledge is produced in a space with two dimensions: one which spans a paradigm-creation-to-paradigm-extension (*R*esearch to *D*evelopment) continuum, and another which spans three discrete categories (*S*cience and Understanding, Tools and *T*echnology, and Societal Use and *B*ehavior). Though these two dimensions seem individually straightforward, we find that in combination they lead to rich and intriguing insights. One of the most important of these insights is: science is not the same as research, and technology is not the same as development; but, rather, science, technology and behavior all proceed along analogous research to development continua.

In Section 3, we introduce the middle layer of the framework. This layer is based on the observation that the character and trajectories of knowledge production in one category (e.g., science) depend strongly on whether or how it interacts with knowledge production in other categories (e.g., technology). We call the various trajectories "streams": 6 are "pure" and do not involve direct interactions between knowledge categories; 12 are "mixed" and do involve direct interactions between knowledge categories. The collection of these streams we call the middle, "knowledge-production streams," layer of the framework.

In Section 4, we explore the top layer of the framework. We call this layer the "knowledge-production systems" layer, because its central feature is that, even as knowledge production in various categories can be described by knowledge-production streams, the knowledge-production streams themselves must link to form a dynamical system of knowledge production. These links are mediated by knowledge itself: knowledge-production streams produce knowledge, and knowledge is the driving force for the knowledge-production streams. This dynamic, made complex by the many types of knowledge and knowledge-production streams, nonetheless lends itself to mathematical treatment via rate equations. We discuss two simple examples of such rate-equation treatments. We also discuss some of the ways in which the links (the effectiveness with which knowledge is a driving force for the knowledge-production streams) can sometimes by weakened by "valleys of death."

Finally, in Section 5 we discuss some limitations of the framework, in particular some of the most important difficulties that would need to be addressed for the framework to make further progress.

2. Knowledge-production space

The foundational layer of the Galileo's stream framework is sketched as the triangular "ST&B–R&D prism" in Fig. 1. It is intended to convey a knowledge-production space with the two dimensions discussed in Section 1, each dimension reflecting an important distinction between different varieties of knowledge production. The first (vertical) dimension has to do with the *stage* of knowledge production—what we generically call the research vs. development distinction. The second (horizontal) dimension has to do with the *category* of knowledge production—what we call the science vs. technology vs. behavior distinction.

2.1. Research and development (R&D)

The vertical dimension of the ST&B–R&D space is based on the notion that knowledge production proceeds

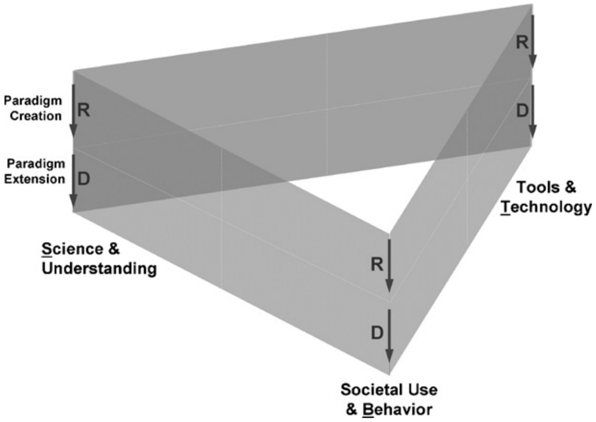

Fig. 1. An ST&B–R&D diagram. The vertical axis represents a continuum of knowledge production increments from Research (paradigm creation) to Development (paradigm extension). The legs of the triangular prism represent three discrete knowledge-production categories: "science and understanding," "tools and technology," and "societal use and behavior."

in stages, and that these stages can be characterized using the concept of "paradigm." In the realm of science, we take this concept to mean something like "accepted examples of actual scientific practice... from which spring particular coherent traditions of scientific research" (Kuhn, 1962, p. 10).

One might say that: we solve puzzles; in solving them we create examples, or paradigms, for solving classes of puzzles; these paradigms grow in generality and influence as they are extended and applied to more and more puzzles.[1] Hence, one can think of knowledge as being produced in stages: "revolutionary" knowledge production, in which new paradigms are created, followed by "normal" knowledge production, in which existing paradigms are fleshed out and extended.

Both paradigm creation and paradigm extension can be challenging and difficult. Revolutionary research is obviously so; but even normal research involves "achieving the anticipated in a new way, and it requires the solution of all sorts of complex instrumental, conceptual, and mathematical puzzles" (Kuhn, 1962, p. 36). However, paradigm creation must surely be considered a larger increment of knowledge than paradigm extension.

The vertical axis of the ST&B–R&D prism shown in Fig. 1 reflects these ideas. If the knowledge produced involves major paradigm shifts and larger increments of knowledge, then it would be represented near the top of the vertical axis; if it involves minor paradigm shifts and smaller increments of knowledge, then it would be represented near the bottom of the vertical axis. Note that we view this distinction as not black and white, but as continuous—whether knowledge production involves paradigm creation or paradigm extension is a question of degree rather than of kind.

Indeed, one might argue that the distinction between paradigm creation and extension is blurry, and one that can really only be made retrospectively. At the time it is being produced, an increment of knowledge is not easy to position, even relatively, along the vertical axis of Fig. 1. Nevertheless, such positioning is considered a central task of virtually all research and development administrators.[2] Hence, we suggest that the paradigm creation and extension distinction *is* important, but that metrics for quantifying the distinction for knowledge production in the present are simply more probabilistic and error-prone than those for knowledge production in the past (Chen et al., 2002).

Moreover, there are indications that the concept of a paradigm is, at the level of the individual, consistent with emerging models of human cognition as nested hierarchies of patterns used in predicting time-sequences of sensory observations (Hawkins, 2004). The leap from a pattern as a mental construct in the mind of an individual to a pattern as a paradigm in the collective knowledge stock of society is a large one, but clearly there are parallels. The counter-intuitive expression attesting to the strength of paradigms, "I'll see it when I believe it," applies just as well to individuals in ordinary life (Yariv, 2002) as to individuals in scientific research communities (Kuhn, 1962).

Of course, we recognize that "paradigm creation" and "paradigm extension" are not the most common of terminologies. Hence, we think it is helpful to capture the distinction between the two with more-commonly used terminology. Perhaps the terminology that best does

---

[1] Note that, whether of a paradigm creation or extension character, what producers of knowledge "do" might be thought of as either "puzzle solving" (following Kuhn, 1962, pp. 36–38) or "problem solving" (following perhaps more conventional usage). A disadvantage of the former terminology is that it seems to trivialize what producers of knowledge do; a disadvantage of the latter terminology is that it seems to imply a use-motivation to the object of study that may or may not actually be present. As we shall see later, we view type of motivation to be a crucial but orthogonal characteristic of knowledge production, and hence opt for the more use-neutral terminology "puzzle solving."

[2] For example, the National Science Foundation, in evaluating research proposals, asks "to what extent does the proposed activity suggest and explore creative and original concepts?" (NSF, 2004, Chapter III); Sandia National Laboratories' Laboratory Directed Research and Development program has as one of its principle objects to "foster creativity and stimulate exploration of forefront science and technology" (DOE, 2001).

so, albeit imperfectly, is "research" and "development." Although there is considerable variation in usage, the word "research" often suggests the possibility, if not likelihood, of creating something new and unexpected; it can perhaps be loosely identified with paradigm creation. The word "development," in contrast, often suggests extending in an expected way something that already exists; it can perhaps be loosely identified with paradigm extension.

Hence, although the word "research" is rich with other connotations, we denote its narrower, paradigm-creation aspect with the symbol R, and associate it with the upper portion of the ST&B–R&D prism. Likewise, although the word "development" is rich with other connotations, we denote its narrower, paradigm-extension aspect with the symbol D, and associate it with the lower portion of the ST&B–R&D prism.

2.2. *Science, technology and behavior (ST&B)*

The second (horizontal) dimension of the ST&B–R&D space is based on our postulate that there are three discrete and qualitatively different categories of knowledge, each with its own community and system of values.

The first category we think of as science and understanding, in which we juxtapose science with understanding to emphasize the notion that the essence of science is to understand. An idealized set of values associated with this community is: universalism, communism, organized skepticism, and disinterestedness (Merton, 1973).[3] These values lead, among other things, to the community of scientists' heavy emphasis on journal articles published in the open literature.

The second category we think of as tools and technology, in which we juxtapose technology with tools to emphasize the notion that the goal of technology is tools that are not ends in themselves, but are means to other ends. Some of the idealized values associated with the community of technologists, such as universalism, are shared with the community of scientists. Some, such as communism and disinterestedness, are not. Ownership of technological knowledge, unlike that of scientific knowledge, is more commonly given to an individual or an institution, rather than to the larger community. These differing values lead, among other things, to the community of technologists' greater emphasis on secrecy (Dasgupta and David, 1994), on patents whose purpose is formal legal ownership, or on technological artifacts themselves (De Solla Price, 1965).

The third category we think of as societal use and behavior, where we juxtapose behavior with societal use to emphasize that it is the behavior of individuals in a society that is paramount, though this behavior is surely moderated by whatever science and technology are (or are not) available. One might think of the values of this community as being those associated with human civilization as a whole, with origins in the complex interactions between human biology and the social and physical environment it is embedded in (Wilson, 1998).

To concretize what we mean by these three communities, it may be helpful to estimate their relative sizes in a modern society such as the U.S. From National Science Foundation figures for 2003 (NSB, 2006, Table 5-11), roughly 1×10^5 science and engineering doctorate holders employed in academia in the U.S. self-identified research as their primary activity. We take this to be a crude estimate of the community of scientists in the U.S., recognizing that it both overcounts (since we mean only to include the community of scientists, not the community of technologists) and undercounts (since the community of scientists is not confined to academia). Also from National Science Foundation figures for 2003 (NSB, 2006, Table 3-1), roughly 5×10^6 people in the U.S. were employed in science and engineering occupations. We take this to be a crude estimate of the community of technologists in the U.S., recognizing again that it both overcounts (since not all of those employed in science and engineering occupations are actually producers of technology knowledge) and undercounts (since the community of technologists is not confined to science and engineering occupations). Finally, in 2003, the population of the U.S. was about 3×10^8 (U.S. Census Bureau, May 2004). We take this to be the size of the overall community of societal use and behavior within which the communities of scientists and technologists reside and with which they interact.

By these rough estimates, then, the community of scientists (1×10^5) is roughly 50× smaller than the community of technologists (5×10^6), which is in turn roughly 60× smaller than the community of the entire nation (3×10^8). It is rather astonishing that the community of technologists can have such a significant impact on the much larger community of societal use with which

---

[3] Universalism means that scientific claims are not accepted or rejected based on personal or social attributes of their protagonist (race, nationality, religion, class, and personal qualities are irrelevant). Communism means that scientific knowledge "belongs" not to the protagonist, but to the larger community. Organized skepticism means that all scientific claims must be subject to testing by the larger community. Disinterestedness means that scientific claims are made without considerations of personal gain.

it interacts; and that the community of scientists can have such a significant impact on the much larger community of technologists with which it interacts. However, like a rudder steering a ship, small increments of knowledge in one category can apparently enable large increments of knowledge in another category.

Our second distinction, then, is based on the notion that these three categories (ST&B) of knowledge are fundamentally distinct. This notion is not without controversy. Particularly in modern times, the processes of creating knowledge in these three categories often rely on a common infrastructure, so it can be difficult to distinguish between them. For example, both the Nobel-prize-winning science of semiconductor heterostructures (Kroemer, 1963; Alferov, 2001) and the technology breakthroughs in high-frequency transistors (Mimura et al., 1980) that enable large-scale societal use of cellular telephony (Larson, 1998) make use of high-purity layered semiconductor materials. But this example illustrates only that there is a complex and beneficial "dancing partner" (De Solla Price, 1965) relationship between the different categories of knowledge (a relationship that we explore further in Section 3), not that the categories are the same. Also, though the outcomes of ST&B knowledge production are very different, some aspects of their production processes can be viewed as similar (hypothesizing, experimenting, re-hypothesizing, etc.). But this only means that there is some commonality to the knowledge-production process in the different categories, not, again, that the categories are the same.

Indeed, that there is a qualitative difference between science and technology has been discussed from the earliest Greek philosophers, who distinguished between *episteme* (knowledge) and *techne* (art) (Layton, 1974), to the latest technology-management textbooks, from which one reads that "science understands nature, and technology manipulates nature" (Betz, 2003, p. 4). Our view is that this difference is deep and fundamental, and one that we see daily as managers and working scientists/technologists. Given scarce resources, a choice must often be made between fabricating a device intended to elucidate *why* a certain behavior is observed, vs. fabricating one intended to show *how* to achieve some performance goal. The choice is often fiercely fought even within a single project team, if the team contains some individuals with a science predisposition and others with a technology predisposition.

That there is a qualitative difference between technology and societal use seems just as deep and fundamental. Technology is the knowledge of *how* to do and make things, separate from an understanding of *what* society might wish to make use of. However, because technology often arises in response to a societal use, and because for many societal uses there exists an obvious enabling technology, it is easy to blur the difference between them. Nevertheless, we argue that there is a difference. Tools and technologies exist, and can be improved, without necessarily being motivated by a societal use. We all know examples of "over tooling" (e.g., striving for the most souped-up automobile engine, the most powerful laser, or the most flexible software code), where improving the tool has taken on a life of its own well beyond known societal uses. Such examples are commonplace in many organizations, including our own Sandia National Laboratories. On the one hand, these "über" tools can seem like sandbox playthings irrelevant to a narrow immediate set of applications. On the other hand, they occasionally provide enormous but unanticipated relevance to a wider set of applications.

The horizontal dimension of the ST&B–R&D prism shown in Fig. 1 reflects the ST&B distinctions discussed above. If the knowledge production is of a science and understanding nature, then it would be placed in the S column. If it is of a tools and technologies nature, then it would be placed in the T column. If it is of a societal use and behavior nature, then it would be placed in the B column. If it is a combination of knowledge production in two categories, then it would be placed in the area between those two columns. And if it is a combination of knowledge production in all three categories, it would be placed in the volume between the three columns.

Note that it is tempting to think of the horizontal ST&B axes as continua, with the character of the knowledge produced gradually changing, e.g., from S to T to B. Our observation, however, is that the three ST&B categories of knowledge production have discontinuously distinct community values and social structures. But just as a person may simultaneously belong to multiple but distinct social communities, knowledge production may simultaneously produce knowledge in the three distinct ST&B categories. Hence, our view is that knowledge produced is best characterized as a weighted combination of discontinuously distinct ST&B categories, rather than as a position along a difficult-to-define ST&B continua.

Also note that these three categories of knowledge roughly map onto emerging notions in cognitive psychology on human memory systems (Tulving, 1985), and onto emerging notions in organizational science on knowledge management (Garud, 1997). Science and understanding maps to semantic, or "know why" knowledge; tools and technology maps to procedural, or "know how" knowledge; societal use and behavior maps to episodic, or "know what" knowledge.

2.3. Independence of R&D and ST&B

The ST&B–R&D distinctions discussed above can be thought of as a visualization tool for mapping knowledge production in terms of two orthogonal (vertical and horizontal) dimensions. An implicit assumption underlying this visualization is that these dimensions, and the distinctions between R&D and ST&B, are mutually independent (or orthogonal). Such independence, if true, has three important consequences.

The first consequence of the independence between R&D and ST&B is that, insofar as all three categories of technical knowledge production follow similar downward paths of paradigm creation and extension, they all have the potential for independent development. None are seen to be necessarily foundational to the others, and none are necessarily elevated to greater or lesser "importance."

The second consequence of the independence between R&D and ST&B is that all three categories of technical knowledge are produced in analogous stages (from paradigm creation to paradigm extension).[4] That this is so in the science category has been discussed extensively (Kuhn, 1962). That this might also be so in the technology category has not been discussed as extensively, but it has been noted that *if* it were so, many observations regarding technology can be explained, including (Dosi, 1982): (a) that technology seems to evolve at times discontinuously, and at other times continuously, and (b) that, once having evolved to a certain point, technologies appear to narrow into a restricted set of possibilities, to the powerful exclusion of others. That this might also be so in the societal-use category has also not been discussed as extensively. However, societal-behavior patterns clearly evolve through similar S-curve-like stages (Rogers and Rogers, 2003) where wary early adoption is followed by large-scale diffusion and dominance.

Indeed, it seems clear that for technology and societal use, just as for science, paradigm extension occupies most efforts, with paradigm creation a much more unusual event. After all, paradigms exist for a reason—they provide an economy to our efforts, allowing us to focus mainly on what has "stood the test of time." Most new ideas in most walks of life turn out not to be useful, and rightly never become elevated to the status of paradigm. When brave new ideas do achieve paradigm status, they will be called different things ("discovery" in science, "invention" in technology, and "revolution" in societal behavior), but they will shape in similar ways how the community perceives and understands reality.

The third consequence of the independence between R&D and ST&B is that, because all three ST&B categories of knowledge production proceed through analogous stages of paradigm creation and extension, none has a monopoly on research or development. In other words, we reject the often-made identification of research with science and understanding, and development with tools and technologies. In our view, science, though sometimes associated with it, does not have a monopoly on research (paradigm creation), and technology, though more often associated with it, does not have a monopoly on development (paradigm extension). Instead, all three knowledge-production categories (science, technology and behavior) encompass both research (paradigm creation) and development (paradigm extension).

2.4. Example of ST&B–R&D diagram

To see the potential utility of the ST&B–R&D diagram presented above, we show, in Fig. 2, an example of a map of the calls-for-proposals for the 13 investment areas of Sandia National Laboratories' 2004 Laboratory Directed Research and Development (LDRD) program. This program is Sandia National Laboratories' main instrument for "stimulating exploration of forefront S&T" (DOE, 2001), and so it is of great interest to understand how its funds are allocated. Note that, for viewing convenience, we have cut the prism between the S and B columns, then unfolded it so that it lays flat around the T column. In this manner, we have reduced the dimensionality of the map to two, at the expense of not being able to represent calls-for-proposals that are combinations of all three knowledge categories. In practice, for a technology-centric laboratory such as Sandia, this does not appear to be a significant limitation.

The map was created in a hybrid process in which manual placement (human judgment) was iteratively refined through comparisons with an automated placement (lexical analysis) process.

The manual placement process was based on a careful reading of the text (~4 paragraphs) associated with each

---

[4] Examples of major and inter-related paradigm shifts in the three categories are: (S) the shift in the first half of the 20th century from classical (deterministic) to quantum (probabilistic) mechanics; (T) the shift in the second half of the 20th century from vacuum tubes to transistors for electrical switching and amplification; (B) the shift in the late 20th and early 21st centuries from off-line to on-line entertainment and learning.

of the calls-for-proposals, followed by manual assignment of two coordinates: a *y* coordinate according to the degree to which the text called for high-risk "out-of-the-box" proposals; an *x* coordinate according to the degree to which the text was aimed at either science, technology, or societal use.

The automated placement process was based on identifying key words and then assigning them *x* and *y* values. For example, the words "insight" and "basic," indicative of science and understanding, were assigned *x* values of 0, while the words "system" and "enterprise," indicative of societal use and behavior, were assigned *x* values of 2. Or, for example, the words "revolution" and "disruptive," indicative of paradigm creation, were assigned *y* values near 1, while the words "mature" and "improve," indicative of paradigm development, were assigned y values near 0. Each keyword was also assigned a variable weight, to reflect an unknown degree of ambiguity. Finally, *x* and *y* coordinates for each call-for-proposal were calculated by counting the number of instances of each keyword, weighting that number by the keyword value and weight, then averaging over all the keywords. The weights were varied to minimize the deviation between the manual and calculated *x* and *y* coordinates.

Based on the deviations, the manual placements, and the keyword value assignments in the automated placement process, were refined. After a few iterations, the manual and automated placements were in good agreement, with the results shown in Fig. 2. Some of the conclusions that can be drawn are listed here:

(1) The most heavily represented ST&B category was tools and technology, with eleven out of the 13 of the investment areas. This is not surprising, given the technology-centricity of Sandia National Laboratories mentioned above.

(2) About half (six) of the investment areas lie on the societal use side of technology. This is also not surprising, given the historical mission orientation of Sandia National Laboratories. Of these, four could perhaps have been anticipated, since they are titled after mission areas: nuclear weapons strategic objectives, non-proliferation and materials control, emerging threats, and energy and critical infrastructure. However, two would have been difficult to anticipate, as they are titled after technology areas and also have strong ties to science: electronics and photonics, and computational and information sciences. It may be an indication of the maturity or immediate-mission-relevance of these technologies that their calls-for-proposals were mapped so closely to societal use. It is also a caution to management that their ties to science, at least in Sandia's LDRD investment portfolio, may be weakening.

(3) About half (six) of the investment areas lie on the science side of technology. This is also not surprising, given the importance of science to advanced technology. Of these, the approximate positioning of four of them could be anticipated: pulsed power, materials science and technology, and engineering sciences, which are positioned towards technology; science and technology strategic objectives, which

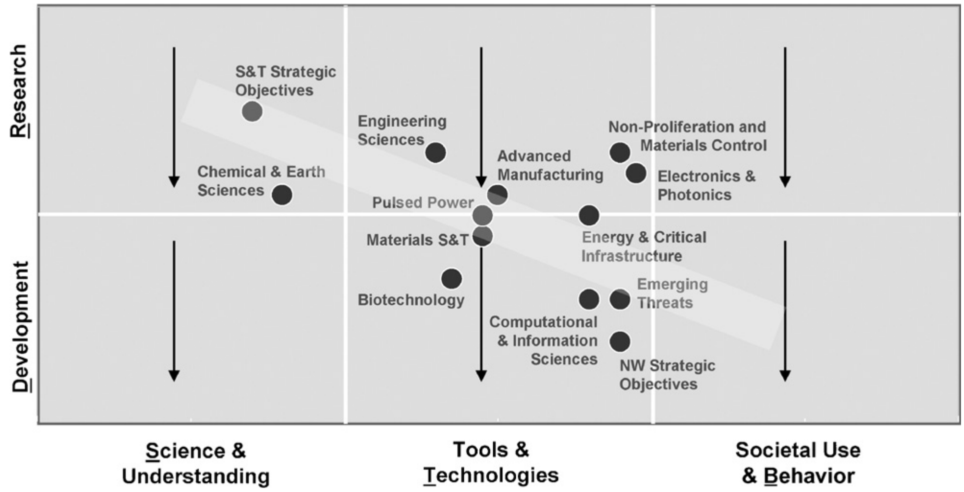

Fig. 2. The 13 investment areas of Sandia National Laboratories' 2004 Laboratory Directed Research and Development (LDRD) program, mapped onto a two-dimensional ST&B–R&D diagram. The translucent white stripe is indicative of a trend that the science-oriented investment areas tend to be more research (paradigm creation) oriented, while the societal-use-oriented investment areas tend to be more development (paradigm extension) oriented.

is positioned towards science. However, two would have been difficult to anticipate. First, biotechnology, which is widely considered to have extremely strong ties to science, is positioned towards technology. This positioning may be an indication that Sandia is a newcomer to this area, and has chosen to enter it by exploiting and emphasizing synergies with its strong existing technology base. Second, chemical and earth sciences, titled partly after a mission area, is positioned towards science. This positioning may be an indication that this area is old to Sandia, and has existing technology strength, but that to tackle challenging new problems (e.g., natural resource utilization and sustainability) stronger ties to science are viewed to be essential.

(4) The two "paradigm maturity" R&D zones are relatively evenly represented. Understandably, the most conservative and "in the box" investment area was nuclear weapons strategic objectives, while the least conservative and "out of the box" investment area was S&T strategic objectives. Indeed, there is a rough diagonal trend (indicated as a translucent white stripe in Fig. 2) in which science orientation correlates to paradigm-creation orientation, and societal-use orientation correlates to paradigm-extension orientation. However, two investment areas deviate from this trend. Nonproliferation and materials control is both societal use and paradigm-creation oriented, an indication that this is a relatively new mission area for Sandia, for which entirely new societal uses may emerge when exposed to existing Sandia technology. Electrons and photonics is also both societal use and paradigm-creation oriented, an indication that existing Sandia technology is being applied to mission areas new to Sandia, with the same possibility of entirely new societal uses.

All of these conclusions are reasonable in hindsight. However, not all the conclusions were obvious in foresight. Hence, such a mapping exercise can provide valuable feedback both to the strategic thinking underlying the calls-for-proposals, as well as to the process by which the strategic thinking is translated into the texts of the calls-for-proposals.

3. Knowledge-production streams

We have discussed, in Section 2, the foundational, "knowledge-production space" layer of the Galileo's stream framework. In this section, we discuss the middle, "knowledge-production streams" layer of the Galileo's stream framework.

One way to think about these streams is as the trajectories of natural sequences of knowledge-production projects. In other words, as one project produces knowledge of a certain character in knowledge-production space, it gives rise to follow-on projects which naturally produce knowledge of a different character.

3.1. Pure streams

In the simplest case, illustrated by the vertical downward-pointing arrows in Fig. 3, one can imagine knowledge production occurring independently in science, technology and societal use. The natural trajectory

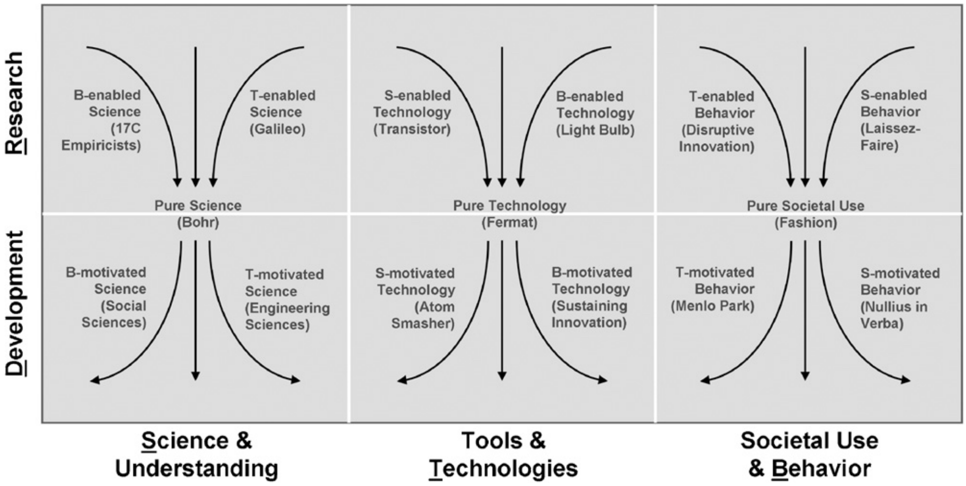

Fig. 3. Knowledge-production "streams," with their exemplars, on an ST&B–R&D diagram. The 6 vertical arrows represent "pure" streams; the 12 curved arcs represent "mixed" streams.

of a sequence of projects then moves vertically down, as paradigms are created in one project, then extended in follow-on projects. Because knowledge production in each of the vertical streams is independent of that in the other streams, we call these "pure" streams.

The far left downward-pointing arrows in Fig. 3 represent the pure-science knowledge-production streams. They can be thought of as the solving of puzzles that advance the frontiers of science, rather than advance or take advantage of a developing technology. They represent science solely for the sake of understanding some aspect of our world. They can perhaps be exemplified by the early 20th century physicist Niels Bohr, who, among other things, created the first quantum-mechanical description of the atom and its energy levels, without apparent regard for the enormous significance that understanding might later have on technology or societal use.

The middle downward-pointing arrows in Fig. 2 represent the pure-technology knowledge-production streams. They can be thought of as the solving of puzzles chosen mainly to advance the frontiers of a technology—rather than to advance or take advantage of either science or a developing societal use. They can perhaps be exemplified by the mid-17th century mathematician Pierre de Fermat, who, among other things, developed new number-theoretic mathematical tools, without apparent regard for the significance these tools might later have on science or societal use.

The far right downward-pointing arrows in Fig. 2 represent the pure-societal-use knowledge-production stream. They can be thought of as the solving of puzzles chosen mainly to advance a societal use—rather than to advance or take advantage of a developing science or technology. They can perhaps be exemplified by clothing fashions, which change from decade to decade, often without direct stimulation from science or technology.

3.2. *Mixed streams*

In the complex but more realistic case, illustrated by the arrowed black arcs in Fig. 3, one can imagine knowledge produced in one ST&B category interacting with and directly influencing production of knowledge in another category. A simple way of accounting for these interactions is to suppose that, to a first approximation, the interactions take the form of the direct application of knowledge produced in one category to production of knowledge in another category. This we assume can be either of two types: knowledge produced in one category *motivating* production of knowledge in another category, or knowledge produced in one category *enabling* production of knowledge in another category.

Because these streams represent knowledge produced in one category interacting with production of knowledge in another category, we call them "mixed" streams. Since there are three ST&B knowledge-production categories, each interacting with two adjacent categories in two different ways, all together there are $3 \times 2 \times 2 = 12$ mixed streams. Four of them are in the S category and involve interactions with the T&B categories; four of them are in the T category and involve interactions with the B&S categories; four of them are in the B category and involve interactions with the T&S categories.

Note that many projects or sequences of projects will not be so idealized, and might have multiple motivations and multiple enablers. In these situations, just as we think of knowledge produced as being a weighted combination of discontinuously different categories of ST&B, we think of an individual project as being a weighted combination of discontinuously different knowledge-production streams. Indeed, one might argue that most projects are such combinations. For simplicity, though, in the remainder of this section we try to give "clean" exemplars of these mixed streams.

Also note that the arcs in Fig. 3 have been drawn in a particular illustrative manner. The reasons will become clearer as we discuss, in the following, the mixed streams in each of the various ST&B categories. However, we do not intend for the exact shapes to be taken too seriously.

3.2.1. *Mixed S&T streams*

The four mixed S&T streams are, clockwise starting in the upper right quadrant: science-enabled technology, science-motivated technology, technology-motivated science, and technology-enabled science.

The science-enabled technology stream can perhaps be exemplified by the semiconductor transistor, a technology whose development in 1947 was enabled by newly developed solid-state and semiconductor physics. This stream is positioned in the upper, research half of the vertical scale to reflect the idea that technologists, when exposed to new scientific insights, are often able to conceive of entirely new technology paradigms. Indeed, the semiconductor transistor could *only* have come about from exploring the ramifications of new science—its design and operation depend on concepts so foreign to everyday experience that its development by chance, without solid-state physics guidance, is nearly unimaginable. The arc points downward and to the right to capture the idea that, as knowledge production focuses increasingly less on exploring the ramifications of the new science, and more on solving particular, narrower technology puzzles, the less likely the creation of new paradigms in technology.

The science-motivated technology stream can perhaps be exemplified by "atom smashers," powerful tools for accelerating particles to extremely high velocities, and for inducing sub-atomic particle interactions.[5] We have positioned this stream in the lower, development half of the vertical scale to capture the idea that technologists, in seeking to develop a tool to benefit science, normally extend existing, rather than create new, technology paradigms. The arc points downward and to the left to capture the idea that, as the choice of technology puzzle becomes increasingly tailored to its relevance to a particular, narrower science question, the less likely it is to extend paradigms in technology in a significant way. The production of each new generation of atom smasher surely required solving a myriad of complex problems, but these were mostly solved through extensions of existing technology paradigms rather than through creating entirely new technology paradigms.

The technology-motivated science stream can perhaps be exemplified by the engineering sciences: the understanding of how engineered tools and technologies "work" and how they might be improved. This stream is positioned in the lower, development half of the vertical scale to capture the idea that scientists, in seeking to explain phenomenon underlying tools and technologies, normally seek to explain them by extending existing, rather than creating new, science paradigms. The arc points downward and to the right to capture the idea that, as the choice of science puzzle becomes increasingly tailored to its relevance to a particular, narrower technology question, the less likely it is to extend paradigms in science in a significant way (see, e.g., Merton, 1973, p. 61). The engineering sciences have certainly played major roles in improving engineered tools and technologies, but they have not often led to the creation of entirely new paradigms in the scientific disciplines they weave together. To pick one example, combustion science has helped improve combustion-engine and jet-propulsion technology, but mostly through the artful combination and extension of existing (rather than creating entirely new) paradigms in its underlying scientific disciplines (fluid dynamics, chemical kinetics, radiation physics).

The technology-enabled science stream can perhaps be exemplified by Galileo Galilei (1564–1642), whose revolutionary astronomical discoveries in the 1610s, including the moons of Jupiter, were enabled by his and others' development of the telescope. This stream is positioned in the upper, research half of the vertical scale to capture the idea that scientists, when using new tools, are often able to make measurements in entirely new domains, and these measurements often enable (or even force) the creation of new science paradigms. The arc points downward and to the left to capture the idea that, as knowledge production focuses increasingly less on exploring the potential of the new tool, and more on solving particular, narrower science puzzles, the less likely the creation of new paradigms in science.

Note that we have named the Galileo's stream framework after this last stream—an indication of the importance with which we view it. Ever since Galileo, at an ever-increasing rate, improvements in observational tools (e.g., microscopes and telescopes of every-increasing resolving power) have been linked to unexpected discoveries and explanations of those discoveries.

Also note that these last two streams differ only by a subtle, but profound, difference in perspective. If we intend to study a tool so that we can ultimately improve its operation (technology-motivated science), we are likely to study it in a narrow range of typical operating environments, a range within which existing scientific paradigms can likely be extended, even if only parametrically. But if we intend to use a tool to explore new phenomena (technology-enabled science) not necessarily related to the operation of the tool itself, then we are likely to expose the tool to a broader range of operating environments in the hope of uncovering unexpected new phenomena. Though Galileo very likely had to study how a telescope works in order to improve it, he is not particularly known for breakthroughs in optical science. But in exposing the tool to the heavens, he made entirely new and unexpected observations.

A subtle change in perspective and motivation can thus be significant. Louis Pasteur studied fermentation, apparently with an initial desire to improve this well-known technology for producing wine. Soon, however, he turned towards exploring fermentation in a wider range of operating environments, particularly those environments which led to normally unwanted products such as vinegar or lactic acid. *Fermentation thus became a tool for studying chemical transformations more broadly, rather than simply a target for process improvement.* And, through use of this tool, he was able to show the existence of biological organisms capable of catalyz-

---

[5] The first atom smasher was a small cyclotron built for roughly $25 in 1931 by Ernest Lawrence (1901–1958) at the dawn of nuclear physics. One of the most recent atom smashers, the Relativistic Heavy Ion Collider, built in 2000 at Brookhaven National Laboratory, is among the largest (two gigantic rings, each 2.4 miles in circumference), most sophisticated (1740 superconducting magnets), and costliest (roughly $650 million) tools of modern science.

ing anaerobic chemical transformation, and to create an entirely new scientific paradigm for chemical transformation.

3.2.2. *Mixed T&B streams*

The four mixed *T&B* streams are, paired diagonally: technology-motivated behavior, behavior-enabled technology, technology-enabled behavior, and behavior-motivated technology. Their positioning and interactions are exactly analogous to those of the mixed S&T streams.

The technology-motivated behavior stream can perhaps be exemplified by Thomas Edison's (1847–1931) famous Menlo Park Laboratory in New Jersey. This laboratory is often said to be the first technology research and development laboratory in history, a behavioral and institutional shift motivated by the need to bring together individuals with different specialties and perspectives to address technology's increasing complexity.

The behavior-enabled technology stream can perhaps be exemplified by the numerous inventions *from* Edison's Menlo Park Laboratory, including the incandescent light bulb. Though Joseph Swan had a decade's head-start, Edison's Menlo Park Laboratory, by bringing together a number of specialties (including glass-blowing, vacuum pumps, electrical theory, dynamos, batteries, and materials and filament making) into an interactive yet focused social setting, caught up quickly and ultimately paved the way for practical electric lighting. Moreover, the incandescent light bulb was only one of a great many technologies enabled and accelerated by Edison's cultural and institutional shift.

The technology-enabled behavior stream can perhaps be exemplified by "disruptive innovation." In this stream, the possibilities of a technology are explored, in advance of (though very likely with the *hope* of) widespread societal use. Every time technology sparks the introduction of a new product or service in advance of a pre-existing pattern of societal use, the risk is high that the product or service will not be accepted, but the potential is also high for creation of an entirely new paradigm in societal use and behavior that disrupts older and competing paradigms. Many, if not most, of the great behavioral shifts in human history fall into this stream: book reading (enabled by the invention of paper and movable-type printing); mechanized labor and transport (enabled by the invention of steam and combustion engines); long-distance communication (enabled by the invention of wired and wireless telephony).

The behavior-motivated technology stream can perhaps be exemplified by "sustaining innovation." In this stream, tools and technologies with an existing pattern of societal use are improved and differentiated in iterative cycles of systematized and routinized (Baumol, 2002) improvement. These improvements are tremendously important, often dwarfing in magnitude the initial improvement enabled by the originating paradigm shift. For example, Robert Noyce (1927–1990) and Jack Kilby's (1923–2005) invention of monolithic semiconductor device integration through planar processing was followed by the Moore's law doubling in semiconductor integrated circuit density roughly every 24 months. This iterative doubling has, from 1971 through 2004, increased transistor count in integrated circuits by 100,000, with no end yet in sight.

Note that these last two streams form a natural pair, but either can precede the other. Disruptive innovation (hard disk drives for personal computing) leads in an obvious way to sustaining innovation (ever-increasing hard-disk-drive densities). However, sustaining innovation can, through incremental improvement of a technology aimed at a particular use (ever-increasing hard-disk-drive densities for personal computing), suddenly reach threshold for a completely different use (hard disk drives for mobile music players), and can thus lead to disruptive innovation (Christensen, 2003). In this manner, "invention is the mother of necessity"[6] leads to "necessity is the mother of invention",[7] which in turn leads to "invention is the mother of necessity," and so on.

3.2.3. *Mixed B&S streams*

The four mixed B&S streams are, paired diagonally: behavior-motivated science, science-enabled behavior, science-motivated behavior, and behavior-enabled science. Their positioning and interactions are exactly analogous to those of the mixed S&T and mixed T&B streams.

The behavior-motivated science stream can perhaps be exemplified by the social sciences, many of which are motivated by a desire to understand how human societies behave and can be improved. For example, much of economic science is motivated by a desire to understand "the nature and causes of," and ultimately to improve, "the wealth of nations" (Smith, 1895). And much of sociology is motivated by a desire to understand the nature of human interactions, and ultimately to alleviate the ills of human societies, such as war, crime, and poverty.

The science-enabled behavior stream can perhaps be exemplified by those instances in which social science

---

[6] Quote attributed to Thorstein Veblen.

[7] Quote attributed to Victor Hugo.

insights have led to changes in how society is in fact organized. For example, “laissez-faire,” the political philosophy of free trade and free markets, was enabled and accelerated by the insights of the early classical economists, particularly Adam Smith (1723–1790) and David Ricardo (1772–1823).

The science-motivated behavior stream can perhaps be exemplified by those instances in which behavioral changes have been made, motivated by improving how science is “done.” For example, the Royal Society, in choosing “nullius in verba” (on the words of no one) as its motto when it was founded in 1660, believed that the quality of scientific knowledge could be improved through reliance on empirical evidence rather than on authority. In doing so, it was of course tailoring and harnessing the growing skepticism directed at religious and other authorities of society as a whole.

The behavior-enabled science stream can perhaps be exemplified by those instances in which cultural changes have indeed enabled an improvement in how science is “done.” The “nullius in verba” behavior discussed above, and propounded even before the Royal Society by Francis Bacon (1561–1626), did indeed unleash new empirical inquiries and entirely new paradigms in science. For example, William Harvey (1578–1657), a close friend of (and private physician to) Francis Bacon, through ingenious experiments on live and dead animals, overturned beliefs dating from Galen (129–200) and established the unity of the arterial and venous circulatory system. This, despite Harvey’s embrace of other aspects of Aristotelian authority and beliefs (such as that the purpose of the blood’s return to the heart is to recover its “perfection”). Other examples from this era include Galileo (discussed already in the context of the technology-enabled science stream) and Anton van Leeuwenhoek (1632–1723), both of whom pioneered the use of experiments rather than authority to establish scientific truth. The overwhelming dominance of this behavioral paradigm, now and in the foreseeable future, is attested to by Richard Feynman’s (1918–1988) dictum “The principle of science, the definition, almost, is the following: *The test of all knowledge is experiment.* Experiment is the *sole judge* of scientific ‘truth”’ (Feynman et al., 1963, p. 1-1).

3.3. Groups of streams

Although the 6 pure and 12 mixed streams are distinct, occupying different areas of the ST&B–R&D diagram, they bear similarities to one another. Through these similarities, an understanding of one stream may improve understanding of other streams. To see these similarities in a systematic way, we group them, as illustrated in Fig. 4, according to the positioning and directionality of their trajectories.

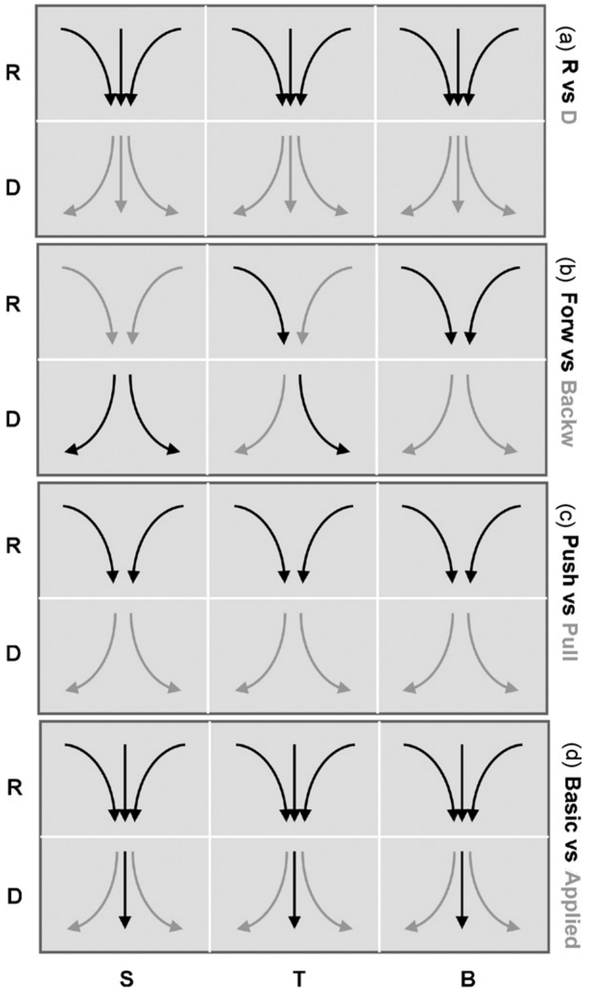

Fig. 4. Similarities and contrasts amongst knowledge-production streams: (a) R vs. D; (b) forward vs. backward; (c) push vs. pull; (d) basic vs. applied.

3.3.1. R vs. D

The first grouping of streams is illustrated in Fig. 4(a). This grouping has to do with whether the trajectories are in the upper, paradigm-creation portion of the ST&B–R&D diagram, or in the lower, paradigm-extension portion. As discussed earlier, paradigm creation and extension can occur in all three knowledge categories. We expect, then, that though the specific knowledge that is created will differ substantially, paradigm creation in science, technology and behavior will share many common characteristics, such as unpredictability, playfulness, hubris, etc. Likewise, paradigm extension in science, technology and behavior will share many common characteristics, such as attention to detail and correctness, emphasis on formality and process, seriousness, etc.

3.3.2. Forward vs. backward

The second grouping of streams is illustrated in Fig. 4(b). This grouping has to do with whether the trajectories are directed “forward” toward what some organizations consider the end point (societal use and behavior), or “backward” toward what some organizations consider the start point (science and understanding).

All the forward streams convey a sense of practicality, a sense of progress being made towards the end point of societal use and behavior. All the backward streams, in contrast, convey a sense of impracticality, a sense in which direct progress is not being made towards the end point of societal use and behavior.

Indeed, it is perhaps this sense of the greater importance of the forward streams that lies at the origin of the “linear” or “pipeline” model. In this model: the pure-science streams produce science; that science occasionally drives a science-enabled technology stream, which merges into the pure-technology streams and produces technology; and that technology occasionally drives a technology-enabled behavior stream, which merges with the pure-societal-use streams and produces societal use and behavior.

We note, though, that over the long term and in a sufficiently large society, the interplay between the various forward and backward streams makes it unlikely that routes to societal use and behavior containing only forward streams will be more effective than those also containing backward streams (see, e.g., Brooks, 1994). It might, for example, be necessary to first develop a tool to solve a key scientific puzzle (a backward stream), before that science could go on to enable a technology (a forward stream) that in turn enables new societal use and behavior (another forward stream). In fact, it is the backward technology-enabled science (Galileo’s) stream that inspired the name for this new framework for knowledge production. That stream is perhaps one of the most important, but often underappreciated, of all the streams.

We also note that, even when considering only forward routes from S to B, there are still two—one that goes directly from S to B and another that goes indirectly from S to T to B. This is easiest to see using the triangular prism in Fig. 1, on which the direct forward route lies entirely on the plane connecting the S and B columns, while the indirect forward route lies both on the plane connecting the S and T columns, and on the plane connecting the T and B columns.

The direct route can be both powerful and simple, as when the understanding that disease originates from germs spurred behaviors aimed at improving household sanitation (Mokyr, 2002, Chapter 5). However, it is unlikely that the direct route is always more powerful than the indirect route. One might even speculate on an historical and perhaps inevitable shift from the shorter to the longer routes accompanied by a growth in technology relative to science and societal use as a fraction of all knowledge production. Perhaps the more advanced science and societal use become, the more difficult further advances become, and the greater the need for ever-more-sophisticated technology.

3.3.3. Push vs. pull

The third grouping of streams is illustrated in Fig. 4(c). This grouping has to do with whether the streams are “motivated” or “enabled.” The motivated streams can be thought of as pull streams, in which knowledge produced in one category is “pulling” on knowledge production in another category. The enabling streams can be thought of as push streams, in which knowledge produced in one category is “pushing” knowledge production in another category. Both push and pull streams are important in an overall system of knowledge production (see, e.g., Schmookler, 1966), but they play different roles and have different characteristics.

The pull streams, because they are generally paradigm extending rather than paradigm creating, are “smoother” streams: the methods they use and the knowledge they produce build on established paradigms within a given knowledge category, hence generate less controversy. Though they are pulled by the hope that the knowledge produced in their category might ultimately lead to knowledge production in another category, failure to do so does not necessarily reflect failure, since knowledge production in that other category is normally viewed as “someone else’s job.”

The push streams, because they are generally paradigm creating rather than paradigm extending, are “choppier” streams: the methods they use and the knowledge they produce build on paradigms that straddle different knowledge categories, hence generate more controversy. If, in using new methods, they only confirm established paradigms, they will likely be met with yawns. If, in using new methods, they overturn established paradigms, they can be easily misunderstood and met with skepticism by a community devoted to those established paradigms.

3.3.4. Basic vs. applied

The fourth grouping of streams is illustrated in Fig. 4(d). This grouping has to do with whether the streams are "basic" or "applied." These are adjectives that we have so far avoided, because they are fraught with ambiguity. However, they are such commonly used adjectives that an attempt must be made to map them onto the Galileo's stream knowledge-production streams.

The most natural such mapping is to think of knowledge production that is "applied" as deriving from projects whose primary motivation is to be of use to knowledge production outside of its category—i.e., all the pull streams. Then, since in common usage "basic" has the sense of being the opposite of "applied," one would then think of knowledge production that is basic in nature as deriving from projects whose primary motivation is to advance knowledge within its category, rather than to be useful outside of its category—i.e., all the pure and push streams.

Note that various philosophical or policy views of the knowledge-production enterprise can be thought of as emphasizing one or the other of these two groups.

One the one hand, the influential views of Thomas Kuhn and Vannevar Bush are most consistent with the basic streams. These streams are unencumbered by a use motivation, hence allow for: (a) optimal choice of puzzles (difficult enough to be challenging, but not so difficult as to be insoluble; Kuhn, 1962); and (b) subsequent *play* with those puzzles ("in order for tool using to develop, it [is] essential to have a long period of optional, pressure-free opportunity for combinatorial activity"; Bruner et al., 1976, p. 38). And, once new knowledge in one category has been produced, it can be harvested through the push streams into other categories, enabling an "endless frontier" of knowledge production (Bush, 1945, p. 15).

On the other hand, the influential views of Gibbons et al. (1994) are most consistent with the applied streams. In this view, the pure and push streams that characterize the traditional "Mode 1" pattern of knowledge production (disciplinary, university-based, autonomous and peer-reviewed) are shifting to the pull streams that characterize a new "Mode 2" pattern of knowledge production (socially distributed, application-oriented, trans-disciplinary, and subject to multiple accountabilities).

Our own view is that both the basic and applied groups of streams are necessary as parts of a functioning *system* of knowledge production (see, e.g., Mowery and Rosenberg, 1979), a view consistent with Donald Stokes' emphasis on all of the quadrants in his "Pasteur's quadrant" framework (Stokes, 1997).

3.3.5. Combinations of groups

In discussing the knowledge production enterprise, it is common to combine groups in various ways. Given the group definitions discussed above, however, only some of the combinations are self-consistent.

An example of a self-consistent combination is one where the adjectives basic and applied are combined with the nouns science, technology and behavior. Basic science can refer to the pure and push science streams (i.e., knowledge production intended to advance science); while applied science can refer to the pull science streams. Basic technology can refer to the pure and push technology streams (i.e., knowledge production intended to advance technology, as discussed by Branscomb, 1997); while applied technology can refer to the pull technology streams. Basic societal use can refer to the pure and push societal-use streams (knowledge production intended to advance societal use); while applied societal use can refer to the pull societal-use streams.

An example of a self-inconsistent combination is one where the adjectives basic and applied are combined with the nouns research and development. Basic research can refer to all the pure and push streams, just as each term individually does. However, applied research cannot refer to any streams, if by applied we mean the pull streams and if by research we mean the pure and push streams.

We suspect there is no resolution to this, other than to acknowledge that the often-used combinations of basic and applied with research are inherently ambiguous and to be avoided. Instead, the combinations of basic and applied with science, technology and societal use, which are not ambiguous, are preferred.

4. Knowledge-production systems

In Section 2 we introduced the R&D and ST&B distinctions that define the foundational layer of the Galileo's stream framework. In Section 3, we introduced the middle layer of the Galileo's stream framework: the idea that knowledge production in various categories can interact to form knowledge-production streams. These streams all have roles to play within an overall *system* of knowledge production. But that system is evidently complex, composed not just of the streams themselves, but of links *between* the streams. This system of linked streams defines the top layer of the Galileo's stream framework.

In this section, we discuss this system of interacting streams, by analyzing two simple subsystems. We concretize the discussion through use of "toy" mathematical models that describe in a simple way the qualitative dynamical features we consider most important. But we

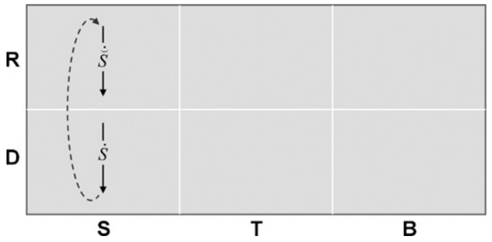

Fig. 5. Knowledge subsystem composed of two interacting pure-science (paradigm creation and paradigm extension) streams.

emphasize that the models are not based on quantitative data, and should be viewed simply as a starting point for more sophisticated and realistic models.

4.1. The “pure-science” subsystem

Let us start with the simple subsystem composed of the two pure-science streams illustrated in Fig. 5: the first representing paradigm creation and the second representing paradigm extension. To represent these streams quantitatively, a simple metaphor might be that of a container partially filled with a stock (or total accumulated amount) of knowledge. As existing paradigms are extended and fleshed out, the stock of knowledge increases and fills more and more of the container. When existing paradigms have become fully fleshed out, “exhausted,” so to speak, the stock of knowledge has filled the container fully. In order for the stock of knowledge to continue to grow, the container itself must expand, through the creation of new paradigms.[8]

Within this metaphor, the first stream increases the size of the container. This size represents a stock of *potential* knowledge, which we label $\breve{S}$, that would result from existing paradigms being fully fleshed out. The second stream fills the container with a stock of *actual* knowledge, which we label S, through extension and fleshing out of existing paradigms. The stock of actual knowledge S is always less than the stock of potential knowledge $\breve{S}$, but the second stream acts to reduce the difference, while the first stream acts to increase the difference.

Let us now suppose that the rates at which the two streams produce knowledge depend linearly[9] on the relative stocks of actual and potential knowledge already produced. A simple coupled pair of knowledge-production rate equations can then be written as

$$\dot{\breve{S}} = \beta_{\breve{S}}\left(\frac{S}{f} - \breve{S}\right), \tag{1}$$

$$\dot{S} = \beta_S(\breve{S} - S), \tag{2}$$

where $\beta_{\breve{S}}$ and β_S are rate constants (with units of, say, year^{-1}). These rate constants could, in turn, be decomposed into the products of knowledge “intensities” (with units, say, of knowledge produced per $) and effort levels (with units, say, of $ year^{-1}), but such a decomposition would be an unnecessary complication for our purpose here.

A qualitative interpretation of Eq. (1) is that the “driving force” for creation of new paradigms is an awareness of residual dissonance that comes as old paradigms are increasingly fleshed out. As the stock of actual knowledge S exceeds some critical fraction f of the stock of potential knowledge $\breve{S}$ (or as the stock of knowledge exceeds a critical fraction of its container), it begins to exert a “pressure” to create improved paradigms and thereby increase $\breve{S}$. That pressure increases as the stock of actual knowledge approaches the stock of potential knowledge. Thus, the paradigm-extension stream provides a “driving force” for the paradigm-creation stream, as indicated by the dashed arc in Fig. 5.

A qualitative interpretation of Eq. (2) is that the “driving force” for fleshing out of existing paradigms is the difference between the unrealized ($\breve{S}$) and realized (S) potential of those paradigms. The rate at which paradigms are fleshed out is highest when paradigms are young and have hardly been fleshed out, and lowest when the paradigms are old and have been fully fleshed out. Thus, the paradigm-creation stream drives the paradigm-extension stream, even as it is being driven by the same paradigm-extension stream, in a coupled positive-feedback loop.

In steady-state (i.e., long times), these two coupled equations imply exponential growth of knowledge proportional to e^{+kt}. In the special situation where the sum of the paradigm-creation and paradigm-extension rate constants is fixed ($\beta = \beta_{\breve{S}} + \beta_S$), corresponding to a situation where the rate constants are proportional to invested

---

[8] This concept is similar to that discussed by Rosenberg (1982, Chapter 7, p. 156) in the context of tools and technology: “It is important to realize that a major technological breakthrough really signals the beginning of a series of new developments of great importance, not their culmination. . . In the most meaningful sense, the development of the transistor or the explosion of the first nuclear device or the first achievement of heavier-than-air flight is really the announcement of a new set of possibilities far more than their attainment.”

[9] One could easily assume other more-than-linear (e.g., combinatorial) or less-than-linear dependences. Here, we assume a linear dependence, as such a dependence leads naturally to the exponential growth consistent with observation.

effort but total invested effort is fixed, the exponential-growth rate constant k can be written as

$$k = \sqrt{\left(\frac{\beta}{2}\right)^2 + (\beta - \beta_S)\beta_S\left(\frac{1-f}{f}\right)} - \frac{\beta}{2}. \tag{3}$$

At either of two extremes, if all effort is placed either solely on paradigm creation ($\beta_S = 0$) or paradigm extension ($\beta_S = \beta$) then the rate constant is zero. In the first case, $\partial S/\partial t = 0$, so S is fixed for all time, while $\breve{S}$ initially increases then saturates at S/f. In the second case, $\partial \breve{S}/\partial t = 0$, so $\breve{S}$ is fixed for all time, while S initially increases then saturates at $\breve{S}$.

Exactly between the two extremes, if effort is placed equally on paradigm creation and extension ($\beta_S = \beta_{\breve{S}} = \beta/2$), then the rate constant has its maximum value of

$$k_{\max} = \frac{\beta}{2}\left(\frac{1}{\sqrt{f}} - 1\right) \tag{4}$$

and the steady-state ratio between the actual and potential stocks of knowledge is

$$\frac{S}{\breve{S}} = \sqrt{f}. \tag{5}$$

Thus, knowledge production in the subsystem is maximized when the efforts applied to the paradigm creation and paradigm extension streams are *balanced*, and minimized when effort is applied *only to one or the other*. This pure-science ($\breve{S}$ and S) subsystem thus illustrates how two linked streams produce knowledge interactively, with the output of the first driving (supplying necessary input to) the second and the output of the second driving (supplying necessary input to) the first.

4.2. The “science-to-technology” subsystem

Let us now consider the subsystem illustrated in Fig. 6. This subsystem, though still illustrative rather than complete, is much more complex than the pure-science subsystem just considered. It has a science half and a technology half, and each of these halves contains paradigm-creation and paradigm-extension streams.

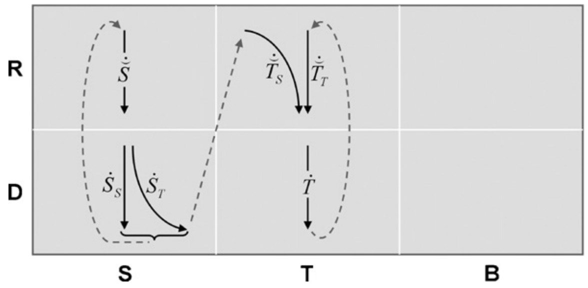

Fig. 6. Knowledge subsystem composed of pure-science and pure-technology streams, and “forward”-linked mixed science and technology streams.

For the science half, as for the pure-science subsystem discussed above, there is a single paradigm-creation stream, though we rewrite it here with a more explicit notation to avoid ambiguity:

$$\dot{\breve{S}} = \beta_{\breve{S}S}\left(\frac{S}{f_{\breve{S}S}} - \breve{S}\right). \tag{6}$$

Unlike for the pure-science subsystem discussed above, however, there are two paradigm-extension streams representing science-motivated science and technology-motivated science, which together sum to an overall stock of science knowledge:

$$S = S_S + S_T. \tag{7a}$$

The rate of increase of these two paradigm-extension streams is given by

$$\dot{S}_S = \beta_{SS}(\breve{S} - S), \tag{7b}$$

$$\dot{S}_T = \beta_{ST}(\breve{S} - S). \tag{7c}$$

For the technology half, the opposite is true. There is only a single paradigm extension stream:

$$\dot{T} = \beta_{TT}(\breve{T} - T). \tag{8}$$

But there are two paradigm-creation streams representing science-enabled technology and technology-enabled technology, which together sum to an overall stock of potential technology knowledge:

$$\breve{T} = \breve{T}_S + \breve{T}_T. \tag{9a}$$

The rate of increase of these two paradigm-creation streams is given by

$$\dot{\breve{T}}_S = \beta_{\breve{T}S}\left(\frac{S_T}{f_{\breve{T}S}} - \breve{T}_S\right), \tag{9b}$$

$$\dot{\breve{T}}_T = \beta_{\breve{T}T}\left(\frac{T}{f_{\breve{T}T}} - \breve{T}_T\right). \tag{9c}$$

The qualitative interpretations of all of these streams are similar to the interpretations, discussed in Section 4.1, of the streams in the pure-science subsystem, with appropriate account taken of which streams are driving which streams, as indicated by the dashed lines and arcs in Fig. 6. For example, the $\partial S_S/\partial t$ science-motivated and $\partial S_T/\partial t$ technology-motivated science streams (Eqs. (7b) and (7c)) in the lower left drive the $\partial \breve{S}/\partial t$ science-enabled science stream (Eq. (6)) in the upper left, and the $\partial S_T/\partial t$ technology-motivated science stream

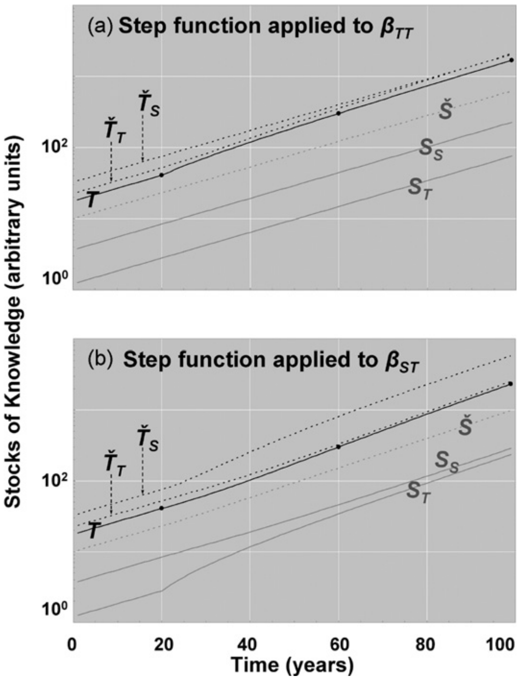

Fig. 7. Numerical simulations of knowledge evolution. Initially, knowledge production is in steady-state exponential growth; at $t = 20$ years, a step function of effort is applied to either (a) β_{TT} or (b) β_{ST}. The time increments for the simulations were 1 year. The initial rate constant parameters were: $\beta_{\breve{S}S} = \beta_{SS} = 0.03\ \text{year}^{-1}$, $\beta_{ST} = 0.01\ \text{year}^{-1}$, $\beta_{\breve{T}S} = 0.05\ \text{year}^{-1}$, $\beta_{\breve{T}T} = \beta_{TT} = 0.02\ \text{year}^{-1}$. The step increases in rate-constant parameters were: $\Delta\beta_{TT} = 0.011\ \text{year}^{-1}$, $\Delta\beta_{ST} = 0.015\ \text{year}^{-1}$. The critical filling fraction parameters were: $f_{\breve{S}S} = 0.20$, $f_{\breve{T}S} = 0.02$, $f_{\breve{T}T} = 0.25$.

(Eq. (7c)) in the lower left drives the $\partial\breve{T}_S/\partial t$ science-enabled technology stream (Eq. (9b)) in the upper right. Hence, they appear as driving forces in these equations.

To visualize the consequences of these equations, we show in Fig. 7 numerical simulations of the evolution of the stocks of knowledge produced by the various streams in this subsystem. There is a wide range of possible rate constants (β's) and threshold fill fractions (f's) that could be chosen for these simulations. Here, we choose a set intended to illustrate a world in which:

- overall knowledge grows at a rate roughly that (2–5% year^{-1}) of the technical (journal articles and patents) literature in recent years,[10] and which can be mimicked in the simulation using rate constants (β's) on the order of 0.03 year^{-1} and threshold fill fractions (f's) on the order of 0.20;
- science-motivated science is more efficient than technology-motivated science at increasing the stock of actual science knowledge (so β_{SS} is larger than β_{ST});
- in the absence of a science-to-technology link via the mixed streams, the growth of science is faster than the growth of technology, so the rate constants for the science streams are (slightly) larger than those for the technology streams ($\beta_{\breve{S}S}$ and β_{SS} are larger than $\beta_{\breve{T}T}$ and β_{TT}), and the threshold fill fractions for the science streams are (slightly) smaller than those for the technology streams ($f_{\breve{S}S}$ is smaller than $f_{\breve{T}T}$);
- in the presence of a science-to-technology link via the mixed streams, growth of science enhances considerably growth of technology, so the threshold fill fraction from science to technology is much smaller than the threshold fill fractions from either science to science or technology to technology ($f_{\breve{T}S}$ is less than both $f_{\breve{T}T}$ and $f_{\breve{S}S}$).

Just as for the pure-science subsystem, the equations for this "science-driving-technology" subsystem have a long-time steady-state in which all knowledge streams grows exponentially with a single rate constant. This steady-state is shown in the first 20 years of the simulations shown in Fig. 7. All stocks of knowledge grow at 4.7% year^{-1}, a rate determined solely by feedback amongst the science streams: in this subsystem we have included science as a driving force for technology, but not technology as a driving force for science. The absolute levels of the technology streams are higher than those of the science streams, because the science streams are driven only by internal feedback, while the technology streams are driven both by internal feedback as well as by feedforward from the science streams. The absolute level of S_S (science-motivated science) is higher than that of S_T (technology-motivated science) because we have assumed $\partial S_S/\partial t$ is more efficient at producing science than $\partial S_T/\partial t$, and has a higher rate constant. The absolute level of $\breve{T}_S$ (science-enabled technology) is higher than that of $\breve{T}_T$ (technology-enabled technology) because we have assumed science is more efficient at enabling new technology paradigms than technology itself is.

In a real-world situation, one can imagine a desire to increase the growth of the stock of specific kinds

---

[10] Over the years 1977–2004, the average growth rate of the technical journal literature, from Thomson Scientific's Science Citation Index for journal articles, has been 4.7% year^{-1}, and the average growth rate of the U.S. patent literature, from the U.S. Patent and Trademark Office's database, has been 2.7% year^{-1} (K.W. Boyack, unpublished).

of knowledge, such as the stock of actual technology knowledge. To achieve this desire, one could imagine increasing the effort applied to the various streams which drive this stock of knowledge. Optimally apportioning effort across the various streams to maximize the increase of technology knowledge would be non-trivial, as the subsystem is a complex dynamical system composed of interacting streams. However, it could readily be accomplished numerically.

To see how the subsystem evolves dynamically in response to changing conditions, in panels (a) and (b) of Fig. 7 we apply a step-function increase in rate constant at $t = 20$ years to two streams. In both cases the magnitude of the increase was such as to cause the stock of technology knowledge to increase by the same amount at $t = 60$ years. Because of differences in how these two streams interact with the overall subsystem, however, the magnitude of the increases are quite different at other times.

In the first case (a), the step-function increase in effort is applied to the technology-motivated technology stream. This is the most obvious stream to apply effort to, since it is the only direct producer of technology knowledge. Indeed, there is an immediate initial increase in $\partial \ln T/\partial t$, the logarithmic rate of production of actual technology knowledge. But because the subsystem is operating under science-limited conditions, $\partial \ln T/\partial t$ eventually reverts to its original steady-state value, although there has been a step jump in the absolute level of technology, T.

In the second case (b), the step-function increase in effort is applied to the technology-motivated science stream. This stream is twice removed from the technology-motivated technology stream, so there are two time lags before the increase in the logarithmic rate of production of technology-motivated science knowledge, $\partial \ln S_T/\partial t$, manifests itself as an increase in the rate of production of actual technology knowledge, $\partial \ln T/\partial t$. Because of these time lags, the short-term (<60 years) increase in technology knowledge is smaller than for case (a). However, the long-term (>60 years) increase is larger, because an increase in the technology-motivated science stream, in increasing the production of science, increases the steady-state exponential growth rate of science, which in turn increases the steady-state exponential growth rate of technology.

We conclude that strategies for optimally apportioning effort across streams to maximize increase of technology knowledge depend strongly on the time scale that is being considered. They also of course depend on the (steady-state or non-steady-state) prehistory of the system, and on the details of the rate constants and threshold fill fractions.

4.3. The "full" system

The full system of knowledge production includes all streams and all possible links between streams. It is clearly complex, even though it is simply composed of analogous streams and links.

The links, in particular, are a critical aspect of the full system. In the toy models discussed in Sections 4.1 and 4.2, the strengths of the links were captured in the "threshold fill fraction" factors. When these are small ($\ll 1$), the stock of actual knowledge in one category represents a large driving force for production of the stock of potential knowledge in the same or another category. When they are large (~ 1), the driving force is small.

In practice, there are many ways in which the threshold fill fractions can be larger than they need, or should, be. This is because the links from paradigm extension to paradigm creation do not occur easily: they are impeded by what might be called "valleys of death," to use a popular term (Branscomb, 2003). Note that there are two kinds of valleys, depending on whether the feedback path is within a knowledge category, or whether it crosses knowledge categories.

4.3.1. Valley of paradigm death

The first type of valley might be called the "valley of paradigm death." It represents a lack of receptiveness of a community to the creation of new paradigms, especially after that same community's old paradigms have become considerably extended and widely diffused.

This lack of receptiveness is reasonable. After all, most new ideas are bad ideas, or at least ideas that are less useful than currently accepted ideas. In science and understanding, history is littered with extraordinary but ultimately false claims; in tools and technologies, history is littered not so much with tried-and-true methods, but with tried-and-untrue ones; in societal use and behavior, history is littered with entrepreneurs who competed unsuccessfully against existing, well-established standards of societal behavior.

Still, there is an obvious negative aspect to this lack of receptiveness. Every once in a while, there is the possibility of what one might call a "wrongful conviction." Normally, communities within a knowledge category follow the rule that new paradigms are guilty (wrong or not useful) until proven innocent (right or useful), particularly in the absence of complete information (Farrell and Saloner, 1985). Following this rule, though, means that the community will normally not invest as much effort to determine the truth or usefulness of a new paradigm as to determine its falseness or lack of usefulness.

In science and understanding, that rule manifests itself in the dictum "extraordinary claims require extraordinary evidence" (Sagan, 1996). In societal use and behavior, that rule manifests itself in the difficulty that an entrepreneur has at obtaining the venture capital (or angel investment) required to progress from idea to realized innovation.

4.3.2. *Valley of application death*

The second type of valley might be called the "valley of applicability death." It represents a lack of receptiveness of a community in one knowledge category to the creation of new paradigms driven by the production of knowledge in another category.

Just as with the valley of paradigm death, this lack of receptiveness is reasonable. After all, the world is complex and, as far as our limited human capabilities go, mostly unpredictable. We may be "pulling"knowledge production in certain directions, hoping that it will be useful to our application, but it is unlikely that it will end up being useful in the way we had hoped. The forward link from S to T, for example, could fail simply because the technology-motivated science that has been developed is not yet adequate to cover the situations of interest to a technology. Thus there would be no possibility of a science-enabled-technology follow-on.

But, just as with the valley of paradigm death, there is an obvious negative aspect to this lack of receptiveness. There could easily be situations where a link could be fruitfully made, but for various reasons the "hand-off is fumbled." Unlike in the first valley, in this second valley the pull and push streams that one hopes to connect reside in different knowledge-production categories (and communities) with different values and cultures. Individuals or organizations working in these different knowledge-production communities may have difficulty communicating with each other. And, worse, they may simply ignore each other—the "safe" path for producers of knowledge in one category surely does not rely on the production of knowledge in another category, especially if it may be difficult for that knowledge to be communicated.

Minimizing these cultural-mismatch-based "fumbled hand-offs" is clearly an important issue for "ST&B transfer" management. Its importance can be seen from the emphasis of the Mode 2 knowledge-production framework on inter-organizational networks and collaborative structures, where "the source of value added lies in the precise form in which the collaboration of groups and the experience and skills of its members takes" (Gibbons et al., 1994, p. 112). Its importance can also be seen from the emphasis of the so-called Triple Helix framework (Leydesdorff, 2000) on the social dynamics between the university, industry and government sectors.

Moreover, in situations where there is a hand-off to be made, and where it can be made smoothly, there is the possibility of a give-and-take between the producer and recipient of knowledge (Amesse and Cohendet, 2001): for not-useful knowledge to be modified or edited into useful knowledge for a given application; or for wrong applications to be modified or edited into right applications for a given piece of knowledge.

For example, for the S to T link, science-enabled technology might lead to an attempt to create a new technology; but enroute a critical unexplained phenomenon might be observed; technology-motivated science then might try to understand the phenomenon, leading to an enhanced technology. In this way, there is a recursive (Brooks, 1994) cycle of science-enabled technology and technology-motivated science that enhances the link between the two. Indeed, if time were an axis pointing out of the page, one might think of these pairs of streams as forming a sort of double helix rising out of the page—similar to that discussed (by Balachandra et al., 2004) in the context of knowledge production in the T&B zone.

Finally, one can consider the S to T *and* T to S links, which gives a spiral consisting of technology-motivated science, science-enabled technology, science-motivated technology, and finally back to technology-motivated science. This might be thought of as a virtuous circle, or "Casimir spiral" (Casimir, 1983, pp. 296–299).

4.4. *System of systems*

Thus far, we have treated each category of knowledge as being served by a single community sharing a common set of paradigms. For most of human history this has not been the case, and even in today's age of globalization is still not the case. Humanity is divided into many communities, often with very little cross-talk between them. This is particularly so for societal use and behavior paradigms, for which customs and norms vary considerably across ethnic groups and nations. But it is also true for science and understanding paradigms, and for tool and technology paradigms.

Hence, it can be useful to think of humanity as being subdivided into *multiple* systems of knowledge production. Each system has its own evolution, determined by its individual streams and by the intra-system links between its streams; but the systems also interact, through inter-system links between streams in one system and streams in another system.[11] Some of these intersystem links will be stronger than others. For example, tools and technologies, implemented in physical artifacts that can easily be transported individually across ethnic and national boundaries, seem more likely to link across systems than societal use and behavior paradigms, which represent a densely interwoven network more resistant to transport of individual behaviors across such boundaries.

Importantly, the strengths of the overall system of systems of knowledge production likely depend very much on these inter-system links. For much the same reason that specialization and division of labor generally are limited by the extent of the market (Smith, 1895), specialization and division of labor in knowledge production are also limited by the extent of the community. In the limit where inter-system links are very weak, then the individual systems will be small, and each will be less productive. In the limit where inter-system links are very strong, then, at least for the knowledge-production streams that are linked, the systems can be thought of as merged, and the whole can be much more productive than sum of the parts.

For example, in the valley of applicability death, just because the knowledge that is being pulled on ends up not being useful to a particular application does not mean that it will not be useful to *some* application—it may serendipitously be extremely useful to other unforeseen applications. The probability that such an unforeseen application will be found increases with the size of the overall pool of applications, and therefore with the size of the overall system.

Moreover, weak inter-system links can cause inferior paradigms to be "locked-into" certain systems. These paradigms could be in any of the categories, but because of the intra-system links across categories, knowledge production in other categories can easily be affected. For example, if in the science and understanding category it is believed that the earth is flat, then in the societal use and behavior category there will not be a desire to sail across the ocean. As a consequence, there can be continued wide gaps amongst nations and geographical regions of the world in knowledge production rates in certain streams, due to inferior paradigms in the systems containing those streams.

Finally, note that individual systems within a system of systems might also compete, rather than cooperate, with each other. The resulting strategies adopted by nations (see, e.g., Nelson, 1993) or industries (see, e.g., Porter, 1990) seeking competitive advantage could be either beneficial or detrimental to knowledge production in the overall system.

5. Caveats

Thus far, we have introduced the Galileo's stream framework, and discussed how, though at this early stage abstract and idealized, it might be useful in quantitative modeling and understanding of knowledge production. Still, we recognize the difficulties of devising a framework to describe a social phenomenon as complex as knowledge production. In this section, then, we mention explicitly three important difficulties, or caveats, associated with the Galileo's stream framework.

5.1. Independence of fundamental quantities

Perhaps the most important caveat is that the fundamental quantities (basically the eighteen knowledge-production streams) in the Galileo's stream framework may have been poorly chosen. In constructing any model, one would like the fundamental quantities to be relatively independent of each other, yet with clear relations between them.

To "pick on" two particular streams in the Galileo's stream framework, one would like, e.g., the technology-motivated science ($\partial S_T/\partial t$) and science-enabled technology ($\partial \breve{T}_S/\partial t$) streams to be relatively independent of each other, yet with the first linking cleanly into the second. However, though it is easy enough to caricature individual researchers in a modern research laboratory as having strong predispositions towards one or the other of these streams, these may only be caricatures—the predispositions and motivations of real researchers are obviously much more complex.

Moreover, predispositions and motivations are not the same as results. To pick on two other streams, even if, e.g., the technology-motivated science ($\partial S_T/\partial t$) and science-enabled science ($\partial S_S/\partial t$) streams were relatively independent of each other, the results of those streams (technology-motivated science (S_T) and science-enabled science (S_S)) might not be. Once having been produced, by whatever stream and by whatever motivation, perhaps science is just science.

All this said, it is not clear what choices of fundamental quantities would be better. Distinguishing between knowledge production by various institutional forms

---

[11] One way of partitioning humanity's knowledge-production enterprise into a "system of systems" would be to consider each nation as having its own "National System of Innovation," and to consider the global system of innovation as mediated by the interactions amongst these national systems (Freeman, 1995).

(e.g., academia, industry, government laboratories) is certainly possible. Distinguishing between knowledge production in various fields (e.g., physics, chemistry, biology) is also possible. We argue, however, that these distinctions are perhaps even more artificial. There are more "process" similarities between, say, science-motivated science in physics and biology than there are between, say, science-motivated science and technology-motivated science in physics. Indeed, we were drawn to the Galileo's stream framework in part because of the common features shared by the various knowledge-production streams regardless of institutional form and field.

Nevertheless, an important challenge for this framework is to understand the manner in which knowledge-production streams differ according to institutional form and field. A related challenge is to understand the similarities and differences between analogous knowledge-production streams (e.g., science-enabled technology and technology-enabled behavior).

5.2. Network vs. rate-equation models

Another important caveat is whether the framework is fine-grained enough to provide a useful "micro" description of knowledge production. In any reductionist approach, one would like to decompose the macro-phenomena into as few simple micro-quantities as possible, but not fewer than necessary.

At one extreme, one can envision network models of knowledge production in which huge numbers of nodes (each representing a chunk of knowledge) and links (each representing an interdependency between two chunks of knowledge) are simultaneously numerically evolved. These kinds of models would, in principle, have the ability to treat arbitrarily complex distributions of micro-scale properties characteristic of either knowledge or knowledge producers, and to simulate their evolution nearly exactly.

At the other extreme, one can envision rate-equation models of knowledge production in which a limited set of key variables interact and evolve. The simplest of these might be that of growth economists, in which knowledge is sometimes lumped into a single "labor-augmenting technology" variable (Jones, 2002). The next simplest of these might be Pasteur's quadrant (Stokes, 1997), in which knowledge is decomposed into four quadrants: Bohr's quadrant (non-use-inspired fundamental understanding), Pasteur's quadrant (use-inspired fundamental understanding), Edison's quadrant (use-inspired non-fundamental understanding), and an un-named quadrant (non-use-inspired non-fundamental understanding). In this progression of models, Galileo's stream, with its eighteen streams, might be viewed as the next–next simplest.

Nevertheless, Galileo's stream is still rate-equation based, is still relatively coarse grained, and may still be inadequate to provide a realistic model for the dynamics of knowledge production. In the long run, it may be necessary to combine rate-equation models with network models. The network models would provide fine-grained micro-scale insight, and the rate equations would provide coarse-grained macro-scale insight into the evolution of mean-field properties of the network (Barabasi et al., 1999). Ideally, some of those mean-field properties will naturally map to the various streams of the Galileo's stream framework. However, it is certainly possible that the mean-field properties end up naturally mapping in some other way, leading to an alternative macro-scale framework for knowledge production.

5.3. Measurements

A final important caveat is whether the framework, at this point a purely theoretical construct, can be compared to measurements, and thus either falsified or provisionally verified. This issue is not unique to the Galileo's stream framework, of course. However, in the absence of a connection to measurement, the framework can never be said to be right or wrong, and its usefulness limited to that of a framework on which to hang anecdotes or qualitative case histories.

We cautiously suggest two possible approaches, neither, unfortunately, without difficulties.

A first approach could be based on the methodology discussed in Section 2.4, in which text-based descriptions of knowledge-production work is used to map that work into one or more knowledge-production streams. The descriptions could be in the form of calls for proposals (as in Section 2.4), the proposals themselves, progress reports, progress evaluations, or final reports. A difficulty with this approach, however, is that, intentionally or unintentionally, researchers are not always accurate in self-assessing their own work.

A second approach could be based on the use of knowledge-discovery-in-databases (KDD) (Fayyad et al., 1996), social/bibliometric indicators (Godin, 2003), or publication and patent citation patterns (Garfield, 1955). It may be possible, e.g., to characterize a publication or patent as belonging to one or more knowledge-production streams through numerical analysis of its backward (bibliometric) or forward citation pattern. A difficulty with this approach, however, is that such citation patterns take time to accumulate, so that comparisons may be limited to historical, not current, data.

Moreover, with both of these approaches, significant additional difficulties are that: they apply mainly to knowledge production that is codified, rather than tacit; and they may not readily apply to the societal use and behavior knowledge-production streams.

6. Summary

In this paper, we have introduced Galileo's stream, a "layered" framework for understanding knowledge production. The framework basically consists of a knowledge-production "space," within which exist eighteen distinct knowledge-production streams, all linked together into a dynamically evolving system. We are cautiously optimistic that the framework, though abstract and highly idealized, will be helpful in quantitative modeling and understanding of knowledge production. However, it remains to be seen whether various difficulties associated with the framework can be surmounted, or whether its usefulness will be limited to that of a framework on which to hang anecdotes or qualitative case histories.

Acknowledgments

We acknowledge valuable discussions with many colleagues at Sandia National Laboratories, including: Cynthia Harvey, Kevin Horn, Keith Ortiz, Hank Westrich, Chuck Meyers, Julie Phillips, Jerry Simmons, Gordon Osbourn, Laura McNamara, Randy Watkins, Gretchen Jordan, Gerry Hays and Pace Vandevender.

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under contract DE-AC04-94AL85000.

References

Alferov, Z.I., 2001. Nobel lecture: the double heterostructure concept and its applications in physics, electronics, and technology. Reviews of Modern Physics 73, 767–782.

Amesse, F., Cohendet, P., 2001. Technology transfer revisited from the perspective of the knowledge-based economy. Research Policy 30, 1459–1478.

Balachandra, R., Goldschmitt, M., Friar, J.H., 2004. The evolution of technology generations and associated markets: a double helix model. IEEE Transactions on Engineering Management 51, 3–12.

Barabasi, A., Albert, R., Jeong, H., 1999. Mean-field theory for scale-free random networks. Physica A 272, 173–187.

Baumol, W.J., 2002. The Free-Market Innovation Machine. Princeton University Press, Princeton.

Betz, F., 2003. Managing Technological Innovation: Competitive Advantage from Change. John Wiley and Sons, Hoboken, New Jersey.

Branscomb, L.M., 1997. The technology politics to technology policy. Issues in Science and Technology 13, 41–48.

Branscomb, L.M., 2003. What's new for technology policy? Issues in Science and Technology 19, 16–18.

Brooks, H., 1994. The relationship between science and technology. Research Policy 23, 477–486.

Bruner, J.S., Jolly, A., Sylva, K. (Eds.), 1976. Play—Its Role in Development and Evolution. Basic Books, New York.

Bush, V., 1945. Science the Endless Frontier. United States Government Printing Office.

Casimir, H., 1983. Haphazard Reality—Half a Century of Science. Harper & Row.

Chen, C.M., Cribbin, T., Macredie, R., Morar, S., 2002. Visualizing and tracking the growth of competing paradigms: two case studies. Journal of the American Society for Information Science and Technology 53, 678–689.

Christensen, C.M., 2003. The Innovator's Solution: Creating and Sustaining Successful Growth. Harvard Business School Press, Boston, Massachusetts.

Dasgupta, P., David, P.A., 1994. Toward a new economics of science. Research Policy 23, 487–521.

De Solla Price, D.J., 1965. Is technology historically independent of science? A study in statistical historiography. Technology & Culture 6, 553–568.

DOE, 2001. DOE Order 413.2A: Laboratory Directed Research and Development. U.S. Department of Energy.

Dosi, G., 1982. Technological paradigms and technological trajectories—a suggested interpretation of the determinants and directions of technical change. Research Policy 11, 147–162.

Farrell, J., Saloner, G., 1985. Standardization, compatibility, and innovation. Rand Journal of Economics 16, 70–83.

Fayyad, U., Piatetskyshapiro, G., Smyth, P., 1996. The KDD process for extracting useful knowledge from volumes of data. Communications of the ACM 39, 27–34.

Feynman, R.P., Leighton, R.B., Sands, M., 1963. The Feynman Lectures on Physics Volume 1: Mainly Mechanics, Radiation, and Heat. Addison-Wesley Publishing Company, Reading, Massachusetts.

Freeman, C., 1995. The 'National System of Innovation' in historical perspective. Cambridge Journal of Economics 19, 5–24.

Freeman, C., Soete, L., 1997. The Economics of Industrial Innovation. The MIT Press, Cambridge, Massachusetts.

Garfield, E., 1955. Citation indexes for science—new dimension in documentation through association of ideas. Science 122, 108–111.

Garud, R., 1997. On the distinction between know-how, know-what and know-why. In: Huff, A., Walsh, J. (Eds.), Advances in Strategic Management. JAI Press, pp. 81–101.

Gibbons, M., Limoges, C., Nowotny, H., Schwartzman, S., Scott, P., Trow, M., 1994. The New Production of Knowledge. SAGE Publications, London.

Godin, B., 2003. The emergence of S&T indicators: why did governments supplement statistics with indicators? Research Policy 32, 679–691.

Hawkins, J., 2004. On Intelligence. Times Books, New York.

Jones, C.I., 2002. Introduction to Economic Growth, 2nd ed. W.W. Norton and Company.

Kroemer, H., 1963. A proposed class of heterojunction lasers. Proceedings of IEEE 51, 1782–1783.

Kuhn, T.S., 1962. The Structure of Scientific Revolutions. The University of Chicago Press, Chicago.

Larson, L.E., 1998. Integrated circuit technology options for RFIC's—present status and future directions. IEEE Journal of Solid-State Circuits 33, 387–399.

Layton, E.T., 1974. Technology as knowledge. Technology & Culture 15, 31–41.

Leydesdorff, L., 2000. The Triple Helix: an evolutionary model of innovations. Research Policy 29, 243–255.

Merton, R.K., 1973. The Sociology of Science: Theoretical and Empirical Investigations. The University of Chicago Press, Chicago.

Mimura, T., Hiyamizu, S., Fujii, T., Nanbu, K., 1980. A new field-effect transistor with selectively doped GaAs–N–AlxGa1−xAs heterojunctions. Japanese Journal of Applied Physics 19, L225–L227.

Mokyr, J., 2002. The Gifts of Athena: Historical Origins of the Knowledge Economy. Princeton University Press, Princeton.

Mowery, D., Rosenberg, N., 1979. Influence of market demand upon innovation—critical-review of some recent empirical-studies. Research Policy 8, 102–153.

Nelson, R.R. (Ed.), 1993. National Innovation Systems: A Comparative Analysis. Oxford University Press, New York.

NSB, 2006. Science and Engineering Indicators 2006. National Science Board.

NSF, 2004. Grant Proposal Guide (NSF 04-23). National Science Foundation.

Porter, M., 1990. The Competitive Advantage of Nations. Free Press, New York.

Rogers, E.M., Rogers, E., 2003. Diffusion of Innovations, 5th ed. Free Press.

Romer, P.M., 1994. The origins of endogenous growth. Journal of Economic Perspectives 8, 3–22.

Rosenberg, N., 1982. Inside the Black Box: Technology and Economics. Cambridge University Press.

Sagan, C., 1996. The Demon-Haunted World: Science as a Candle in the Dark. Random House.

Schmookler, J., 1966. Invention and Economic Growth. Harvard University Press, Cambridge, MA.

Smith, A., 1895. An Inquiry into the Nature and Causes of the Wealth of Nations. G. Routledge.

Stokes, D.E., 1997. Pasteur's Quadrant. Brookings Institute Press, Washington, DC.

Tang, S., 2005. Knowledge as Production Factor: Towards a Unified Theory of Economic Growth. Preprint.

Tulving, E., 1985. How many memory-systems are there. American Psychologist 40, 385–398.

U.S. Census Bureau, May 2004. Annual Estimates of the Population for the United States and States, and for Puerto Rico: April 1, 2000 to July 1, 2003 [http://www.census.gov/popest/states/NST-EST2003-ann-est.html].

Wilson, E.O., 1998. Consilience: The Unity of Knowledge. Knopf, New York, NY.

Yariv, L., 2002. I'll See It When I Believe It—A Simple Model of Cognitive Consistency. Cowles Foundation Discussion Paper No. 1352.