

- ◆ Introduces a way of working with drug development teams that relies on writing as a key development activity
- ◆ Argues that structured writing can help team members engage in substantive conflict and reach consensus on difficult issues

Knowledge Management and Pharmaceutical Development Teams: Using Writing to Guide Science

STEPHEN A. BERNHARDT AND GEORGE A. McCULLEY

INTRODUCTION

In our consulting work at McCulley/Cuppan LLC, much of what we do involves helping drug development teams work collaboratively to prepare new drug application documents for the global market. We coach teams in implementing document development processes that place a strong emphasis on initial conceptual planning, early document prototyping, early drafting, and public reviews. We fight the tendency of scientists to first do the science and *then* write the documentation. We try to integrate doing the science with writing, so the writing can guide the science and the teamwork.

Successful development of a new drug requires close, well-articulated collaboration across a range of disciplines, from synthetic chemistry to medicine, with major contributions from chemical engineering, analytical chemistry, toxicology, medicinal chemistry, pharmacology, and other well-defined scientific disciplines. The new drug application (NDA) should present a well-articulated, consistently interpreted set of key messages and issues. A strong application will coherently present to the regulatory authorities the best thinking and argumentation the team has to offer for approval of the drug, with the desired label. (A label contains the prescribing instructions for physicians and is packaged with each drug dispensed. The drug's label separates it from competitors.) In a cross-functional team environment, it is a challenge to get the team working together to align the data with the strongest possible arguments for drug approval. The approved label provides marketing leverage and can determine the success or failure of the drug product.

This article describes the beginning step (which we call the *seed document*) in a systematic document development process for new drug teams as they work collaboratively to produce the necessary approval documentation. We demonstrate why we work with teams to capture their knowledge in seed documents, and then discuss how seed doc-

uments lead to document prototypes and drafts. These same seed documents play a role in document reviews and ultimately in assessing a dossier.

All examples used in this discussion are real, though the company identities, the exact nature of the drugs and their qualities, and any data represented have been altered to protect proprietary information. The pharmaceutical industry is driven by proprietary knowledge, and concerns for information security run high.

THE PROBLEM

Within a cross-functional development team, data, information, and strategic knowledge are widely dispersed across individuals from different functions within the organization. A major challenge is finding ways for teams to work together to consolidate what they know and present it consistently across a set of documents.

The development teams must produce NDA dossiers of considerable length and complexity. A filing might run from 200,000 to 600,000 pages and contain hundreds of individual documents. The filing brings together complex data sets, representing chemical development, animal research, and clinical trials in humans. (For background on drug development, see Rubin 1984; for specifics on clinical studies, see Spilker 1986; for general background on New Drug Application documentation, see Bonk 1999.) The writing of the actual documents is spread across the cross-functional teams, with many individuals contributing text modules or entire draft documents to the filing. Other authors work in the background, not as primary team members but as supporting personnel within line functions. So the situation is one where a multidisciplinary team is co-authoring a very long and complex document set. They must call on highly distributed data and knowledge,

Manuscript received 15 July 1999; revised 31 August 1999; accepted September 1999.

using authors from across a large, bureaucratic organization. The organization itself likely retains many of its functional alignments in addition to newly formed cross-functional teams.

Although there are some trained medical writers within most pharmaceutical companies, many of the authors are scientists and technicians first, and authors second, and they frequently have had no formal training in scientific writing and rhetoric. Drug development projects take a long time (typically 6 to 12 years), and inevitably, many team members and authors have not previously experienced a filing or written a particular genre of filing document before, so they do not necessarily have strong experience as lead authors.

Cross-functional teams are relatively recent innovations, and ways of working are evolving. For team members, thinking of writing as a collaborative process, as opposed to an individual effort, does not come naturally ("I'll show you my draft report when it is finished"). Their models for writing tend to be scientific articles, and they tend to favor what they perceive to be a scientifically objective presentation of facts, as opposed to seeing their task as analyzing the data and developing the available means of persuasion within logical, well supported arguments.

In fact, drug development data seldom (if ever) speaks for itself, so teams must develop complex and subtle arguments on the basis of often equivocal data. For example, some HIV drugs have been initially approved, at least partially, on the basis of whether they increase the numbers of a particular type of white blood cell despite the fact that the connection between longevity and wellness and the counts of this white blood cell is tenuous. So any arguments about efficacy in this indication are founded on

Their models for writing tend to be scientific articles, and they tend to favor what they perceive to be a scientifically objective presentation of facts, as opposed to seeing their task as analyzing the data and developing the available means of persuasion within logical, well supported arguments.

somewhat indirect measures of disease remediation. Such situations lead to complex and subtle arguments that require careful articulation. Indeed, the practice of using so-called "surrogate markers" as the primary data supporting approval of new drugs is currently receiving critical attention from the regulatory authorities; they would naturally like to see actual, direct measures of disease remediation.

The situations in which NDA texts are externally reviewed are also complex. Drug regulatory agencies exist in most countries worldwide, setting varying and complex regulatory guidelines that determine the content and order of the documentation. Although efforts have been made to "harmonize" regulatory guidelines, these efforts are still far from complete, and it is a daunting task to write a single set of documents that meets the varying and constantly changing requirements around the world. The reviewers themselves will be experts, but the potential range of their expertise is very wide, and the team can never presume that the reviewing authority will have a specific knowledge set matched to the authors' expertise. There is a recognition that authorities in various global markets will review filings differently, but it is very difficult to tailor documents to individual markets and to the localized attitudes toward disease and toward the risks/benefits of specific treatments.

For example, in France health authorities expect very aggressive drug therapy in the treatment of Parkinson's disease, while in the United Kingdom the expectation is much more conservative. All reviewers can be presumed to read "against the text," looking for problems, gaps in the data, weak efficacy data, unexplained or unrecognized patterns of adverse events, or difficulties meeting specifications for purity or consistency in the drug product. So the review situation is somewhat unpredictable, and reviewers can be expected to be antagonistic.

Finally, the problems confronting drug development team authors are compounded by the nature of scientific investigation, which is rooted in inductive logic, with the results of one study leading to another. Scientists who author the documents tend to follow this inductive approach as they write their reports, a stance reflected in the use of inductive rhetoric and organization in the reports. The documents tend to resemble detective novels, building from one revealed piece of evidence to another, often with large leaps in the logic, to grand conclusions about the relative safety and effectiveness of the drug. The most important information is frequently not in prominent places in the filing hierarchy, but spread throughout as it arises in the context of discussion. The documents often do not meet the needs of reviewers, who are less concerned about the journey the scientists took and much more concerned with rendering an opinion on the end product: Is it safe? Does it work? Does it meet a real need? Is it worth the

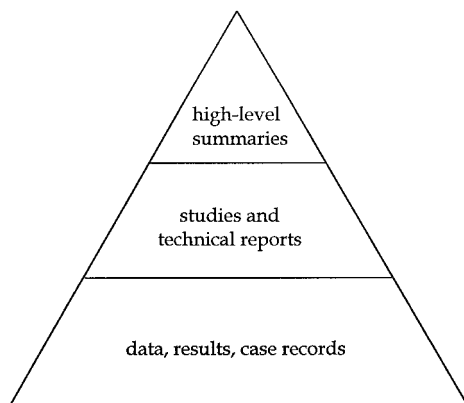


Figure 1 The deductive pyramidal structure of a New Drug Application (NDA).

money?

The superstructure of the dossier reverses the inductive nature of science by creating a pyramid organization, with data sets and case records at the base, supporting reports at the middle level, and top-level summaries at the apex, as in Figure 1.

This pyramidal structure means that there must be internal consistency up and down through the document hierarchy, and that many key messages and critical issues must be represented redundantly, but at varying levels of specificity, at high, middle, and low levels of the dossier. Reviewers are sometimes presumed to devote first attention to the high-level summary documents, but we know from interviews with health authority readers (on the basis of consulting interventions at the U.S. Food and Drug Administration) that reviewers are unpredictable in their approaches to the documentation, with some going to the data first to gain an unbiased view, some going to the key study reports, and others going to top level summaries to get the big picture. The dossier must be accessible at all levels, must be internally consistent, and must convey the most important messages in emphatic positions.

These complex situational demands pose complex challenges to the development team.

THE SEED DOCUMENT WITHIN THE LARGER DOCUMENT DEVELOPMENT PROCESS

We view writing as an essential tool for productive scientific and team-based collaboration, and we use a seed document as the first step in a collaborative writing process. The seed document enables teams of drug developers to use writing as a constructive activity throughout the drug development process.

In the initial compilation of the seed document, the team brainstorms to capture all the issues associated with a

project. They examine the competitor's drug labels (the approved marketing authorization) to see how the drug under development compares to the competition and to see what issues emerge regarding the efficacy or safety of the new drug. They think out loud, area by area, giving voice to all the anticipated development challenges and problems they face, capturing each in writing and representing it as an issue. At team meetings following the initial construction of the seed document, as new issues arise, they are recorded in the seed document. Often, only some of the columns for any row are filled in—an issue may be identified well before there is any logical response developed to address it.

Table 1 presents a typical portion of a seed document for a project to develop a novel treatment for dental plaque (all details altered)—a lozenge that actively disintegrates dental plaque. The **Issue** column sets the challenging question, the troublesome bit. The **Response** column captures in a declarative statement the position the team will argue: their current position. The **Rationale** column captures their logical argumentation or support for this interpretation: their warrant for making the claim as represented in the Response. The **Support** column points toward studies, data, publications, or other evidence that can be marshaled or that must be completed to support the response and rationale.

The seed document is *issue-focused* for a purpose—we want teams to concentrate their energies on the most difficult development challenges. We also want those challenges or problems to be posted in writing in full view of the team, so the combined talents and specialized expertise of the various team members can be brought to bear on formulating and testing appropriate responses to the tricky issues of development. We spend a lot of time nailing down the issues, getting them represented in uncontested language, so that everyone agrees on the statement of the issues.

In addition to identifying the issues, the seed document is a place to put tentative interpretations of the science—passing theories—that give way to more fully articulated and defensible interpretations as the patterns in the data become clearer. The seed document also allows the team to explore conflicts in how the members define the key issues and what must be done to address those issues. The seed document is a physical stage, if you will, where debates over the science can be acted out.

Notice that the cell contents in the seed document tend to span functional areas of development. Issues frequently point toward chemistry (ion deposition) as well as to preclinical studies (*in vitro* studies of ion deposition on enamels). Some of the issues begin with chemistry or preclinical studies, such as the concentrations of heavy metals, but move out into manufacturing and controls (how to manu-

TABLE 1. PORTION OF A SEED DOCUMENT FOR A BIOLOGICALLY ACTIVE PLAQUE-REMOVING LOZENGE

Issue	Response	Rationale	Support
At least some users are likely to chew the lozenge and cause abrasive action on teeth enamel.	The lozenge will be abrasive, since the raw material is inherently abrasive, so an acceptable level of abrasivity must be defined and controlled via chemical manufacturing.	The raw material shape and size properties can be altered and controlled through milling to eliminate grittiness. Some abrasivity is characteristic of any lozenge, including breath mints and candy, so some level of abrasivity is acceptable.	<i>In vitro</i> : Midwest State University study assessed abrasivity of gum on teeth—need to recover this study for details. We need to define specific levels of abrasion or grit.
What are the consumers actually swallowing when using the lozenge? How much can safely be ingested? Do we need a recommended daily dose limit?	With regard to the active ingredient, consumers swallow nothing of any concern from a toxicological perspective. From a legal perspective, we are covered up to our recommended daily limit on the active substance.	Tox assessment shows active substance is safe up to 1500 mg/kg/day, a level approximately 1300 times the anticipated clinical dose under current usage recommendations. The clinical program will address use of the product.	<i>In vitro</i> release data will determine—TBD/Clinical data. The tox program evaluated the active in a 13-week study of oral administration at <3500 mg/kg/day. Due May 2000.
Does the lozenge deliver solution ions to the front teeth? Is the activation with existing plaque layers well distributed or concentrated around the back molars? Should subjects and users be instructed to move the lozenge around to cover all teeth?	The lozenge delivers solution ions to and between the front teeth. It is not necessary for people to intentionally move the lozenge around to different mouth areas.	We don't have a strong argument until we produce some <i>in vitro</i> data. In theory, saliva will distribute active substance throughout the oral cavity, with ion disposition reaching comparable levels across all teeth after approximately 10 minutes.	<i>In vitro</i> data supporting release of ions (fluoride data) due March 2000. May need an <i>in vitro</i> trial to test for deposition distribution via saliva in the absence of intentional action.
Have we met all regulatory requirements for registering the lozenge as a food? What about concentrations of heavy metals?	We will attempt to meet all regulatory requirements for registering the lozenge as a food, though it may not be practical in some markets. In some EU countries, the requirement for metals is as low as 1 ppm, which may not be reachable. We will need to decide if we are not going into certain countries because we cannot meet their guidelines for heavy metals.	Requirements are clearly defined, and in some markets requirements to limit heavy metals, and especially lead, are quite stringent. The active has some lead impurities (in part because it is produced in factories that make ceramics). The levels of lead may go down as factory production is devoted solely to the lozenge (the equipment gets cleaner). But there are also trace quantities of lead and other metals in excipients used in the manufacture of the lozenge.	Regulatory needs to gather country-by-country data on heavy metal specs for food and cosmetics for all intended markets. We will need support for market decisions by June 2000.

facture the product and keep the levels of heavy metals below specifications) and toward marketing and business decisions (what country-by-country limits are established for lead content in food items). Data on the chemical manufacturing and controls (such as starting lead content in the active drug substance and excipients supplied to the company from various chemical suppliers) ties into broad marketing decisions far downstream—as in the anticipated inability to meet requirements in some markets and movement toward a decision not to pursue licensing of the product in certain countries. The knowledge captured in this seed document helps pull together the upstream and downstream development issues, from suppliers upstream to markets downstream.

Notice, too, that there are several kinds of language in the cells. Some of the language consists of well-formed, complete sentences, portable to draft reports. Other language is tentative—the best statement of position at a certain stage of development. Other cell contents are simple placeholders—what must be done to actually have a well-stated position or a strong argument. The seed document contains evolving language that will eventually be formulated in ways that declare issues resolved, state clear positions, point toward supporting data, and argue persuasively for approval. For this product, filing is still many months away and issues are still being raised, debated, and resolved.

The seed document grows as more is known about the drug, often reaching lengths of 100-plus pages. This should not be surprising, given the size of the entire new drug application. The complexity of the seed document determines appropriate media for handling the tabled information. Sometimes the seed documents are simple word processor tables; other times they are created in a spreadsheet and can be sorted and filtered. The Excel® spreadsheet from which Table 1 was clipped had additional columns to flag critical path issues and to assign issues to organizational functions. In addition to allowing the user to sort on key terms, the seed document can be filtered to show only those issues that cross over, for example, manufacturing and regulatory. We have also worked with a company to program the seed document as a relational database, with links from the various issues to threaded e-mail conversations, to the emerging dossier with

its drafts at various stages, and to the supporting studies. With a Web-site instantiation, there is the potential to make the seed document part of the company intranet, linking the development team to contract research organizations and to the health authorities themselves in a fully participatory development initiative.

Our use of seed documents has its roots in fundamental assumptions of writing theory—theory familiar to the readers of this journal. We use writing as essentially a conceptualizing, *constructive* activity, as opposed to a simple *recording* activity. That is, as writers think their way through troublesome or tricky issues, and as they work out their responses on paper, they are doing the conceptual work of drug development. *Writing* is a way of *doing* science and a way of *working* as a team. It is impossible to predict the outcome of the heuristic activity of writing in formulating responses and rationales. Writing here is a thinking tool, a way to reason in a team environment and to test definitions of issues alongside potential or alternative responses. In the words of Michael Schrage, “It takes shared space to create shared understandings” (1995, p. 94). The seed document is a first space in a collaborative writing process that produces a shared understanding among team members.

John Seely Brown, Chief Scientist and Director of the Xerox PARC laboratory, discusses a relevant example from their work with a pharmaceutical company, Syntex, during a project named Express, which studied drug development teams in action:

... one of the most interesting lessons we've learned from the Express project so far is just how long it takes to create a shared understanding among the members of such project teams—a common language, sense of purpose, and definition of goals. This is similar to the experience of many interfunctional teams that end up reproducing inside the team the same conflicting perspectives the teams were designed to overcome in the first place. We believe the persistence of such misunderstandings may be a serious drag on product development. (1997, p. 172)

Brown goes on to detail the ways that information technologies might accelerate mutual understandings within work groups, as does Schrage. Undoubtedly, such technologies have a role, and we have discussed technologies for global drug development teams elsewhere in detail (Bernhardt 1999, Bernhardt 1995). Here, we would stress that the solutions do not need to be high tech and that, in fact, simply getting the shared team thinking in one place in writing is a simple way of moving toward consensus. We cannot overstress the crucial role of getting it in writing. It is only as one is forced to nail down the issue and response in specific terms that the team really knows what their position on an issue is and where they still need more supporting data.

Sometimes the seed documents are simple word processor tables; other times they are created in a spreadsheet and can be sorted and filtered.

Does the writing of the seed document actually produce the collaboration? No, not really. In our experience, people generally are willing to work together. The problem is that collaboration on cross-functional teams is often not very efficient. People sit in team meetings literally for days, come back in a month and discuss the same issues that they had “resolved” the previous month. What writing does is make the collaboration concrete, scrutable, and lasting—the foundation for future necessary documentation.

SEED DOCUMENTS AS THE INITIAL STAGE OF DOCUMENT DEVELOPMENT

The seed document is a first step with the team, a step that begins a process of document development that relies on initial prototyping, early drafting, and public team-centered reviews. The seed document and subsequent steps in the collaborative writing process are relatively simple tools with substantial implications for coordinating cross-functional teams in a product development environment.

Placing text in a table cell is a way of beginning the writing process, of getting words on paper, without facing a blank sheet of paper and the accompanying procrastination or writer’s block. The process structures the first phase of document development, offering writers a good starting point for actually sitting down and drafting a report. Furthermore, the activity of creating a seed document is a team process, so that the project moves along with shared understanding of the critical issues and the team’s responses to those issues. This approach makes writing a social activity from the beginning. It stands in contrast to typical ways of working, when other team members review documents only at late stages of development. In our consulting interventions, team members participate in invention activities (brainstorming the issues, working out tentative responses, identifying sources of support or gaps in support) and review activities (reviewing the seed document and crafting the language that captures the issues and responses) at early stages, when conceptual and strategic formulations are most malleable.

Further, working initially in a conceptual space, such as the seed document, alleviates the concern or preoccupation writers normally have with formulating complete text on paper or screen. That is, the seed document is a loose, tabular, logical structure, rather than a tight, linear, sequential structure. Cells can be placeholders, reminders, partially formed ideas that will not be lost and can be more fully developed later in the writing process. The evolving text within the cells can gain definition over time, until the seed document contains fully formed sentences and fleshed-out logic that can be exported to a prototype report. The process is efficient because the language that is crafted within the cells can eventually be pasted into reports as reusable text modules. The process also controls

1.3 Important preclinical development issues

Pharmacokinetic Issues: There is an apparent delay in T_{max} (time to maximum plasma concentration) in dogs, but this is believed to be an artifact of two methods of measuring, with the HPLC being insensitive to that fraction of the active which is avidly bound to proteins. In contrast, the UV assay showed much less variability and a more rapid T_{max} , as would be expected with a bolus infusion.

This is not surprising since the HPLC assay has not been fully developed, and the limit of detection of the UV assay is currently better, perhaps because the UV recognizes even tightly bound substance. When we write the report, we need to compare the limits of detection of the two assays to explain the different kinetic findings. This same logic needs to carry over to the Chemistry, Manufacturing, and Controls analytic methods validation report.

Unresolved issue: The potential problem in humans is that the protein binding of the active is likely to be higher than in the dog, and the limit of detection of the UV assay probably cannot be increased any further. We could potentially have no viable assay for the active in humans. The question is whether we should fund to continue the development of the HPLC assay now or wait and see if we really need it when we go into humans. If we wait, it would slow our development by 2 months.

Figure 2 Sample document prototype section, with issues from seed document in place.

quality, as the document database—the text modules in the seed document cells—helps ensure that different authors pose issues identically and offer consistent responses, no matter where reviewers encounter those issues throughout the complex NDA dossier.

Notice that the seed document is not an outline: it does not map specifically to a document structure. While it contains the “seeds” that can populate a prototype document, the seed document itself is a tabular representation (not a report structure) of the current thinking of the team on all critical issues. Notice, too, that the seed document is not a project management or document development plan. Mapping out tasks and timelines is a separate and equally necessary procedure. While the seed document is not a project plan, it does frequently alert the project planner to activities that are time sensitive or that are in the critical development path.

Issues from the seed document can be mapped directly into various documents, populating them with appropriate content. We encourage teams to move from the seed document to report prototypes—early maps of the report that identify key messages and issues, determine a page count, identify graphics and tables, and make decisions about what is covered in a text, what is covered in an appendix, and what links are established between one report and another. Such prototype reports are more dynamic than traditional outlines—they show structure and strategy, content and approach.

Figure 2 presents a sample of a first prototype of one small section of a preclinical (animal studies) expert

report, a high level document that sits atop the NDA pyramid.

Notice in Figure 2 that the initial prototype is not yet an outline, nor is it a draft report, though it shares some characteristics with both. Rather, it is a proto-form to hold the issues and responses generated in the seed document. The author is not yet at the stage to identify all topics discussed, in point-by-point order. Nor can the author offer an authoritative response to the issues. At this stage, the prototype is populated with the issues that demand attention in a given section of a given report. It contains language that may hold through the report generation, and it focuses the author on the important issues to be discussed in the report. So it provides very good guidance to authors on the direction of report development (and development of the science, too).

This use of the seed document to spawn the prototype filing documents has several consequences. Most importantly, the seed document reminds authors to concentrate on the most important drug development issues. In a work environment where reports can be very long (a complete clinical trial report might have 600 pages, including reams of attached data and record forms), authors need help in representing the key development issues and making sure they are prominently represented and logically argued. The seed document puts issues in plain view—as they are copied from the seed document to the report prototype, they claim important, emphatic positions in the report hierarchy. Moving from seed documents to prototypes helps authors escape becoming trapped in a narrative of development, retelling the history of one study after another. The process forces a focus on development *issues* as opposed to *history*.

The process also constrains the natural impulse of many authors to subordinate troublesome issues, whether they are development issues that can still be fixed or issues related to claims being made by the company in its registration dossier that must be logically supported. It is tempting in long reports to bury the troublesome parts and hope reviewers will not catch them. But this strategy risks alienating the reviewer, who is on a mission to find the weak spots in the filing. The strategy is risky because the attempt to hide problems with a drug is an act of bad faith, dam-

aging to the implicit partnership between the pharmaceutical company and the health authorities, who share the goal of bringing safe, effective, and necessary drugs to market. Reviewing authorities can be particularly hard on companies if they feel the data does not warrant the sought-after approval—they can request additional (and expensive studies), they can restrict the marketing authorization, and they can stop a filing in its tracks.

Forming a prototype around the key issues prevents the author from simply arriving at issues as the text unfolds topically. Instead, the author foregrounds the key issues at the beginning of the report, and then emphatically addresses those issues with fully developed argumentative responses. The arguments respond to the issues and show how the supporting data and response warrant the company position. This allows authors to take control of the issues, and the reports allow reviewers to read with an eye toward evaluating the arguments, warrants, and quality of the data, rather than reading to discover flaws or gaps in the filing.

The use of prototypes shares much in common with Killingsworth's approach, as presented in his technical communication textbook, *Information in action: A guide to technical communication*. Killingsworth advocates the use of CORE documents, document plans that map out purpose, audience, content, strategy, visuals, page count, and so on, before students begin to write. Such prototypes or CORE documents force authors to think at high, strategic levels before attempting to outline or draft details of a report.

The seed document with prototyping encourages certain important changes in how individuals manage their writing processes. Most importantly, the seed document makes writing a more team-centered, public affair. Because the seed document works at a high, abstract, conceptual level, it encourages collaboration in the early stages, before a document is a draft. This practice stands in sharp contrast to usual ways of working, where nobody sees a draft document, or discusses its content or rhetorical strategies, until late stages, when the draft is circulated in full form for review.

In practice, there is typically substantial resistance to putting work in front of the team before it is "finished." It is typical for an author to invest a large amount of energy, ordinarily at late stages in the project, creating a fully-fledged draft document for line function or team review. These late-stage drafts are created under intense pressures from the imminent filing deadline. Such late-stage drafting and review lead to poor quality documents because of hurried reviews and resistance to making major, strategic changes to completed drafts. We actively counter this resistance in our consulting interventions, attempting to make writing a more open, public, team-centered activity,

Moving from seed documents to prototypes helps authors escape becoming trapped in a narrative of development, retelling the history of one study after another.

Good writing is usually the product of a methodical work process rather than a last-minute burst of effort.

with early creation of a seed document, early prototypes and drafts, and early, high-level reviews.

Early reviews can focus on the seed document and early prototypes, to evaluate whether the important issues have been addressed, and if so, whether the strongest arguments have been put forward. The author can resolve “tough” issues early on. The result is fewer debates once all the data is in, with review sessions that go more smoothly because some of the tough disagreements have already been negotiated. This is early, high-level, collaborative strategic review, as opposed to late-stage, adversarial review. In our model, late stage reviews can confirm that the documents capture the best arguments based on the final data, conform to regulatory requirements, and are consistent with other documents in the dossier.

When writers and other contributors are able to spread their work over a longer period of time, they normally produce better work. Those of us who teach writing know the consequences of not seeing a draft until late in the writing process—it becomes increasingly difficult to suggest major strategic revisions once the author has committed the document to paper and solidified the approach. Good writing is usually the product of a methodical work process rather than a last-minute burst of effort. While the period of time used to construct the report may grow, the amount of time (number of hours worked) will actually fall and fewer rewrites will be necessary. These benefits accrue because decisions have been made throughout the process, issues surface while they can be easily addressed, and communication has been achieved throughout development.

THE RESULTS WHEN GOOD TEAM PROCESSES ARE NOT FOLLOWED

We have many examples of the bad consequences that follow when good team processes for cross-functional communication do not exist. Making the document development process collaborative and public opens communication channels that encourage the team to exchange critical information.

Not long ago we were asked to review some documentation concerning a drug that lowers blood pressure. We quickly noticed that the drug did not adequately lower blood pressure if it was taken only once a day; it had to be taken

twice daily. The problem is that blood pressure drugs have a narrow therapeutic window (the range in the amount of drug that is effective and still safe). Failure to take one of two doses a day would result in inadequate control of blood pressure. Worse, accidentally taking a third dose per day could be disastrous. Therefore, all modern blood pressure lowering drugs are formulated for once a day dosing. People can generally remember to take a drug once per day, but make many errors with twice or three times per day dosing. Such dosing would be impossible to sell.

How did a drug with such an obvious flaw proceed in development to the point where it was actually filed? Simply put, individual members of the project team didn't communicate. The drug formulation people (chemists) didn't know until it was too late that there was any concern about the formulation (which they would have solved if they had known). The clinical development members of this team (the people who conduct the studies in humans) knew that the length of effective blood pressure control was “iffy” but believed they were stuck with the existing formulation. Clinical people generally care about efficacy and safety data, assuming that marketing and business can take care of selling a particular dosing regimen. Preclinical and clinical pharmacology people knew that the drug was quickly eliminated and would not maintain a sufficient concentration in the plasma without twice daily dosing, but they tend to think about how the drug behaves once inside animals and people, not about how it should be formulated and dosed. Business and marketing were not even on the development team—they would inherit the work following approval, when the product was “thrown over the wall.” They would have seen the issue immediately and insisted on changing the formulation or forgetting the project, because they have to sell drugs—and attendant dosing regimens—all the time.

The collaborative writing process that we practice, based on developing a seed document, can first identify and then keep unresolved issues in tension, encouraging competing interpretations or positions in the light as the team sorts out the issues and reaches consensus. The seed document process would have specifically asked whether there were any issues associated with frequency and amount of dosing. This is a “rhetorical topic” of development, a question that is always asked. The idea of the seed document is to get the alternative viewpoints, the conflicts, and the competing needs into the open and then debate the issues until resolution.

We have some empirical evidence that teamwork is improved by such substantive conflict. Burnett (1997), for example, distinguishes substantive conflict (arguments over competing ideas, approaches, explanations) as inherently fruitful to a writing team's productivity and the quality of their documentation. Burnett carefully distinguishes

Do project teams already have to work effectively together to use our process?

such conflict from other, non-productive personal conflict (where people don't like each other and act out) or procedural conflict (where a team does not share clear rules for interacting and achieving its goals). Her controlled study (1993) shows quite clear gains in the quality of writing when teams experience relatively high levels of substantive conflict. Note that Burnett's notion of conflict does not mean vituperative argument or heated debate, though it can sometimes take that form. Burnett argues that the best teams collaborate most effectively when they purposefully create a period of deliberations where alternative positions are articulated and considered by the team. We would argue that both the writing and the science benefit from processes that intentionally bring issues to the full team's consideration.

We would argue further that there is a qualitative difference in process and outcome when such substantive conflict, or exploration of alternatives, takes place in writing as well as in talking. People can talk a lot in meetings, and they can believe that they share an understanding of an issue or an interpretation of data. But as Schrage argues, "In an oral conversation, the words have a soap bubble quality: they float around, evoke some comment, and then pop and disappear" (1995 p. 69). Schrage argues for WYSIWIS ("whizzy whiz") tools, where "what you see is what I see" (p. 68). The seed document is a WYSIWIS tool, one that records the collaborative work of the team and that gives all team members voice and access.

Do project teams already have to work effectively together to use our process? No, but it sure helps. We have successfully worked with project teams who couldn't agree about much of anything. These teams, fortunately, are the exception rather than the rule. In these instances, the amount of work required of everyone is probably doubled. We have observed, however, that making someone's ideas public by putting them in writing and sharing them tends to either wither resistance or well up agreement. So our process tends to encourage teams to concentrate on best alternatives and to mitigate unproductive disharmony.

CONSOLIDATING THE DISTRIBUTED KNOWLEDGE OF TEAMS

The processes we advocate are not aimed solely at producing strong documents, but in facilitating the knowledge work of teams. Team management tends to be centered in a project leader, who must look out for all project areas,

identify activities on the critical path, make resource adjustments as necessary, and keep the company informed about what issues might delay registration or even stop a project altogether. Project managers have the tricky duty of pulling together team members from across the line management hierarchy of the organization to work collaboratively on a development team. Members of the teams have twin allegiances—to the project and to their line function—and work under twin systems of supervision and reward. The knowledge of the team is distributed across team members and line functions, but the project manager and the team as a whole need consolidated knowledge.

The seed document can be a place where the allegiances of the team members can be brought together and knowledge consolidated. There is a powerful synergy in getting highly educated and highly experienced experts from different disciplines together to identify and address project issues. The team can put pressure on its members to be forthright and to identify for the team those problems that arise in the different functional areas. This kind of openness is not immediately comfortable for all team members, but teams exert powerful self-disciplining moves on recalcitrant or reluctant members, once standards have been established. This openness can also be maintained throughout the collaborative writing process.

An example of a successful intervention can show the value of our recommended seed document process to managing team knowledge. We worked with a team developing a new arthritis agent, one of the new class of COX-2 inhibitors that have recently offered an alternative to aspirin and other non-steroidal anti-inflammatory agents. In the early going, as the team formed up and started mapping its issues to a seed document, the chemists, in particular, attempted to stay on the sidelines. Their stated position was that "We don't have any issues that the team needs to be aware of" and that "If any issues do arise, we will take care of them." It took pressure from the project leader and team members to pry open the chemists and to start the process of identifying potential issues. That other team members, representing other functions, had already started recording the major issues in their areas paved the way for the chemists.

With some prodding, the production chemists, who make the drug in large batches, suggested there might be some difficulty producing enough drug in formulation (in actual pill form). The project team had been forced by management to shorten the development time from 3 years to 2. Thus, the original plans had to be accelerated for producing enough drug for animal and clinical trials. When the issue of drug supply was placed in the seed document, it generated a lot of discussion, with attendant issues, since planning for animal and clinical studies, as well as drug launch following approval, would all be affected if there

were problems simply producing enough drug. This particular drug demanded a complicated synthesis, and a single batch took many weeks to produce. There were no shortcuts.

Identifying the issues got the team to start making calculations of drug substance (the chemical itself) and drug product (the pill form) for all the studies that the various team members were planning. There were many surprises among the various members, since no one had anticipated and computed overall figures. Nobody had the big picture; each team member or group knew only what they would need for their studies and assumed there was no supply problem (and some had assumed others were aware of their unstated needs). The production chemists themselves were shocked, as they had never been briefed on the quantities needed for all the various studies, and they followed up with discussions with their managers to detail the project needs.

The company had some 30 other projects at various stages of development, with some big drugs close to launch, and as the chemists talked with their line managers, it became clear that the company did not have sufficient capacity to produce the necessary drug substance because of prior commitments to other projects. This came as a surprise to the team chemists, who knew only their own projects. Pretty soon, planning escalated to budgeting for a new factory.

What happened in this case was a result of a team using a certain tool, the seed document, to begin mapping out issues in a public, collaborative fashion. Putting the production issues into the seed document led to team members discussing and revealing their own needs for drug supply and led, in turn, to discussions with management about overall production capacity. What began with assurances that there were “no issues” on the chemistry side quickly escalated to management decisions involving tens of millions of dollars for building a new chemical production facility.

It is important to note that *no one* had the big picture, not the team leader, not the line management within production chemistry, and certainly not the chemistry team members themselves. It was the process of information sharing stimulated by constructing the seed document that started the process of knowledge sharing, contingency planning, and resource allocation.

The knowledge was in the group, but it was distributed rather than consolidated. Perhaps if the right people had had the right conversations, the production problem would have been anticipated. More likely, the production problem would have emerged as a crisis much later in time, when the opportunity to plan had disappeared, and the project would have been delayed.

Creating the seed document provided both the stimulus

and the physical location for the team to capture and manage what they knew. *Knowledge management* is admittedly a consulting buzzword, which many view with suspicion as nebulous and ill defined. (See Davenport and Prusak 1998 for a helpful map of the field.) In this instance, we think knowledge management is an appropriate way to think about a collaborative, systematic writing process, using a specific tool, to help manage team knowledge in a structured fashion.

GETTING THE SEED DOCUMENT PROCESS TO WORK IN A TEAM ENVIRONMENT

It would be nice to think that a useful process implements itself, but in reality, the use of new tools requires a cultural change and is likely to encounter various forms of resistance. Garvin describes this difficulty well in his article “Building a learning organization”:

Most training programs focus primarily on problem-solving techniques, using exercises and practical examples. These tools are relatively straightforward and easily communicated; the necessary mind-set, however, is more difficult to establish. Accuracy and precision are essential for learning. Employees must therefore become more disciplined in their thinking and more attentive to details. They must continually ask “How do we know that’s true?”, recognizing that close enough is not good enough if real learning is to take place. They must push beyond obvious symptoms to assess underlying causes, often collecting evidence when conventional wisdom says it is unnecessary. Otherwise, the organization will remain a prisoner of “gut facts” and sloppy reasoning, and learning will be stifled. (1998, p. 54)

What Garvin describes is true in our experience: it is a simple matter to conduct a training course and introduce tools such as seed documents, or to demonstrate how to prototype, draft, and review a report. It is much harder to change work practices so that new ways of working are integrated with the everyday work of science.

We witnessed an incident recently that underscores this point. Earlier in the project, the development team had drafted a statement of the indication for a new cancer therapy that would concentrate on refractory tumors—those that had been treated with chemotherapy, or radiation, or other means (sometimes all other means), but that had proved resistant or recurrent. Thus,

To make the seed document a tool that works in practice, teams need to commit to new ways of working.

the team was positioning the therapy as second-line (behind other more conventional therapies) or even third-line (as last resort). In the process of developing their seed document, however, and as a result of much discussion and argument, it was decided to take the aggressive, risky position, and to broaden the indication to include first-line therapy, so that prescribing oncologists could choose to use the new cancer therapy as an initial alternative to chemotherapy or radiation (a much more inclusive label indication and therefore a bigger market). As this argument played out, the team recognized a need for supporting data, which, in this case, was thin. But they still felt they could make a case and gain approval, and so began building documents and the filing around the broader indication. These decisions—the issue of first-line indication, the rationale, the intended statement on the drug label, the supporting data—were recorded in the seed document, which was available (theoretically) to all team members through project folders on a Lotus Notes Domino server.

In actuality, some team members relied on paper and memory, rather than the Notes groupware, and so ended up working from earlier decisions. Since the team's work was captured electronically in the seed document, those who worked from paper handouts from earlier meetings were out of date. In this case, the unfortunate outcome was that the team set up meetings with the FDA and with European health authorities and sent meeting documents based on the more narrow indication. Some team members were shocked that others had not kept up electronically; others were dismayed that such important decisions were not copied and distributed through conventional paper distribution channels. All agreed that simply having an understanding of a tool does not solve the problems of implementing it in the practice of everyday work.

To make the seed document a tool that works in practice, teams need to commit to new ways of working. In a perfect world, team members would exhibit the following attitudes and practices:

- ◆ Be willing to work cross-functionally, to understand others' areas and issues, and to see the value in bringing together people with differing expertise
- ◆ Be willing to be forthright about the development issues—be willing to put in writing the most troublesome and challenging development tasks
- ◆ Be willing to put partially formed responses, raw prototypes, and embarrassingly "drafty" drafts in front of other team members, including those from other line functions, for strategic review
- ◆ Be willing to return from time to time to the seed document and evolving report prototypes and drafts to see that all the issues are captured and that the

responses and support are lined up in the most effective arguments

- ◆ Be willing to work with the seed document, to know what represents current information, and to go online to accomplish work if that is where the current versions reside
- ◆ Be willing to write reports that put main messages and issues in prominent positions, and that directly address the most troubling areas of development prominently and with the best available means of persuasion

In reality, however, on every team there are members who do not endorse and practice all these actions. If they did, it would make our work easier and the process of developing seed documents smoother. But even in an imperfect world, where team members disagree about best practices, the seed document process can foster collaboration. While some resist the process passively or actively, others see the advantages of working around a seed document, change the ways they work, and improve the output of the team.

In our experience, bringing about cultural changes in the work of teams is a substantial challenge and not quickly accomplished. It helps if there is somebody with good understanding of document development processes and team knowledge management who can *drive* the process.

Davenport and Prusak (1998) suggest that corporate librarians are well situated to undertake key roles in knowledge management, an interesting proposition that reflects the traditional roles of librarians and expands on their expertise in very natural ways. Librarians know how to catalog and maintain archives, and they know their way around electronic tools for sorting and sharing information. Librarians, however, are not typically the people with the best knowledge of document production processes. They may not know how to work with authors to create the most effective documents. Davenport and Prusak concentrate on capturing organizational knowledge and making it accessible, a traditional approach to knowledge management that is somewhat different from knowledge work with specific outcomes

... even in an imperfect world,
where team members disagree
about best practices,
the seed document process
can foster collaboration.

(as in our case, where the filing documents represent tangible outcomes of shared knowledge).

As external consultants, we have frequently been the drivers of the process, taking authority for managing the seed document, keeping it up to date and in front of the team, suggesting to the team when it would be useful to review and update the seed document, explaining the rationale for the process and coaching recalcitrant individuals, bringing the seed document back into view as various project documents are first prototyped, and making the key issues and responses the grounding criteria for document review.

Developing seed documents and prototypes is a face-to-face, collaborative process; it is truly an exercise in collaboration and team building. Trained facilitators are often key to the success of prototyping sessions for a development team that has never been introduced to the process. Facilitators and consultants are expensive, but the cost can be justified in an industry such as pharmaceuticals, where time to market is critical and where development costs can be recouped rapidly if the drug generates strong sales.

Another approach is to center authority for knowledge management on the team member who represents Regulatory Affairs. These are the individuals whose task is to communicate with regulatory authorities through correspondence and meetings, and to serve as advisor and team coordinator for the filing, making sure that the dossier reflects current regulatory guidelines and that it satisfies all requirements. These individuals are naturally positioned to work across functions. They also tend to have specialized knowledge that is particularly relevant to filing documents, since it is their job to know the regulations governing what and how drug information is reported to authorities. They also have very specific knowledge of the generic conventions of specific filing documents (for example, how a U.S. summary differs from an E.U. summary, and what is required of the documents in each case).

One company we worked closely with over several years found the consulting work we performed sufficiently valuable that they created their own department, essentially building team facilitation into the corporate structure. Other companies have decided that medical writers (the technical writers of the industry) are well equipped to take on the role of knowledge manager for the team. This makes a lot of sense, since these writers have strong writing skills and an appreciation for what good documents look like. The downside here is that medical writers are frequently accustomed to working as individual writers, who are handed information by a particular research and development function and who then write a report and circulate it for review. This kind of isolation of writing from the team development and management of knowledge is neither an inherently rich practice nor a strong power base.

A better practice, but one which involves changing the fundamental roles of writers in organizations, is to make a place on the cross-functional development team for the writer, so that the writer is involved in strategic discussions and brings document development expertise to the group, while taking on the new role of managing the team knowledge as reflected in the seed document. One company with which we worked placed a medical writer on each team as the documentation expert. This person had expanded duties—not just as report author but as the person who “owns” the seed document, who facilitates prototyping sessions, who leads the team in electronic knowledge sharing and documentation practices, and who helps the team keep track of what they know and what they will argue. Placing this person on the team is recognition that “document science” or “knowledge management” in a cross-functional team can add value, just as having a statistician or a business economist on the team can add value. Doing so makes a lot of sense, since the whole effort of the team is directed toward pulling together knowledge in documents to support a marketing approval—a lot of drug development *is* document development. Such a change parallels those we have seen in the computer documentation industry, where the writer becomes a member of the hardware or software development team.

CONCLUSION

We have reviewed in this paper an approach to a consulting intervention intended to improve the quality of the science and the documentation that drug development teams produce. Our emphasis has been on one step of the collaborative writing process that we practice in the pharmaceutical industry. As our primary illustrative example, we have focused on the seed document, a written, visual, shared repository of team thinking, meant to become the locus for idea development and issue resolution over time as a team works toward its goals. However, all other steps in this process, from document prototyping through team-based reviews, provide the same locus. We see all steps in the process as both thinking tools—for shaping the consensus shared among team members on the project—as well as writing tools. The goal is to support strong science and to produce high quality documents that get the expected results. ♦

REFERENCES

- Bernhardt, S. A. 1995. “Technology-driven documentation in the pharmaceutical industry.” *Journal of computer systems documentation*. 19, no. 4:13–18.
- Bernhardt, S. A. 1999. “Using technology to support global drug-development teams.” In *Exploring the rhetoric of professional communication: An agenda for teachers and researchers*, ed.

- C. R. Lovitt and D. Goswami. Amityville, NY: Baywood Publishing Company, pp. 55–80.
- Bonk, R. J. 1999. *Medical writing in drug development: A practical guide for pharmaceutical research*. Binghamton, NY: Haworth Press.
- Brown, J. S. 1997. "Research that reinvents the corporation." In *Harvard business review on knowledge management*. Boston, MA: Harvard Business Review Press, pp. 153–180. First published January-February 1991 (Reprint 91101).
- Burnett, R. E. 1993. "Conflict in collaborative decision making." In *Professional communication: The social perspective*, ed. N. R. Blyler and C. Thralls. Newbury Park, CA: Sage Publications, pp. 144–162.
- Burnett, R. E. 1997. "Collaboration in workplace communication." Chapter 5 in *Technical communication*, 4th ed. Belmont, CA: Wadsworth Publishing Co., pp. 85–114.
- Davenport, T., and L. Prusak. 1998. *Working knowledge: How organizations manage what they know*. Boston, MA: Harvard Business School Press.
- Garvin, D. 1998. "Building a learning organization." In *Harvard business review on knowledge management*. Boston, MA: Harvard Business Review Press, pp. 47–80. First published July-August 1997 (Reprint 97402).
- Killingsworth, M. J. 1996. *Information in action: A guide to technical communication*. Boston, MA: Allyn and Bacon.
- Leonard, D., and S. Straus. 1997. "Putting your company's whole brain to work." In *Harvard business review on knowledge management*. Boston, MA: Harvard Business Review Press, pp. 109–136. First published July-August 1997 (Reprint 97407).
- Rubin, A., ed. 1984. *New drugs: Discovery and development*. New York, NY: Marcel Dekker, Inc.
- Schrage, M. 1995. *No more teams: Mastering the dynamics of creative collaboration*. New York, NY: Currency/Doubleday.
- Spilker, B. 1986. *Guide to clinical studies and developing protocols*. New York, NY: Raven Press.