THE DISUNITY OF KNOWLEDGE WORK

Jeremy Aarons

School of Information Management and Systems,
Monash University, Australia
jeremy.aarons@sims.monash.edu.au

Session I-5

Abstract

This paper applies recent discussions about the implications of a disunified view of the sciences to the practical challenge of understanding and managing knowledge work in modern organizational settings. The main claim of this paper is that a disunified view will provide a rich and powerful analysis of knowledge work, and thus help provide the necessary guidance for the support of knowledge work in organizational environments. In particular, the disunified view can provide a methodology for analysing knowledge work in a complex organizational setting, and can provide guidance on how to manage organizational change when it involves a fundamental shift in the nature of knowledge work.

Keywords: disunity of science, knowledge work, managing organizational change, agile development.

The Disunity of Knowledge Work

Jeremy Aarons[a]

[a] School of Information Management and Systems
Monash University, Australia
jeremy.aarons@sims.monash.edu.au

Abstract

This paper applies recent discussions about the implications of a disunified view of the sciences to the practical challenge of understanding and managing knowledge work in modern organizational settings. The main claim of this paper is that a disunified view will provide a rich and powerful analysis of knowledge work, and thus help provide the necessary guidance for the support of knowledge work in organizational environments. In particular, the disunified view can provide a methodology for analysing knowledge work in a complex organizational setting, and can provide guidance on how to manage organizational change when it involves a fundamental shift in the nature of knowledge work.

Keywords: disunity of science, knowledge work, managing organizational change, agile development.

1 Introduction

The task of understanding and managing complex organizational environments is clearly a formidable challenge, and an increasingly important one in today's rapidly changing world. Indeed Drucker (1999) identifies the task of improving knowledge worker productivity as the biggest challenge for management in the 21st century. This is especially the case when the organizational setting involves highly trained and skilled workers, engaged in a highly technical enterprise that incorporates high-level technical expertise in a collaborative environment, utilising complex technological systems.

Yet despite the widespread acknowledgement of the increasing importance of supporting such knowledge work in organizational management, organizations are still struggling to approach this challenge in a systematic and rigorous way. A case in point is the Australian Bureau of Meteorology. In this case, although the imperative to evolve work practices and develop new weather forecasting systems is generally recognised throughout the organization, organizational change has so far being driven largely in a piecemeal and ad hoc type of way. Admittedly there has been considerable progress in framing a general model of an improved system (Linger et al, 2001), and a few working subsystems are being developed which conform to this general model (Kelly et al, 2003). But whether this approach is sustainable, and will achieve its long term goals, is not at all clear. For as yet the pathways for developing and implementing these new systems fully across the organization are not entirely clear. As a result the long term future of the redevelopment process is uncertain.

This paper takes up the challenge of finding a suitable methodology for the study of complex organizational systems of knowledge work. Such a methodology is clearly needed for guidance on how to manage organizational change, especially when it involves a fundamental shift in the nature of knowledge work. The approach recommended here will be achieved by applying some recent developments in the philosophy of science to the specific domain of knowledge management (KM). Here an appeal is made to recent work in the philosophy of science which explores the methodological and metaphysical implications of the disunity of the sciences (Dupré 1993; Galison 1997; Clarke, 1998; Cartwright 1999). This work challenges many of the long held assumptions about the way scientists work, the development of scientific theory, and the very nature of science itself. As a result, a new view of science emerges, one that is not only a more accurate picture of the scientific enterprise, but also has clear methodological implications for the practice of science. In particular, it allows for a plethora of methodologies and theoretical frameworks, which can make use of a whole range of technologies, modelling methods, and computational techniques.

Following from this, the broad claim of this paper is that just as a disunified view of the sciences has methodological implications for the practice of science, a disunified view of knowledge work has significant practical implications for managing knowledge work in complex organizational environments.

At this point one may wonder why the philosophy of science is at all relevant to the domain of knowledge management and the complex practical challenge of improving knowledge worker productivity. Why should the methods used to study the physical world be at all appropriate for the study of organizational systems or knowledge workers? The response here is that although the endeavours of science and KM may seem fairly distinct, they do in fact have much in common. In particular, both enterprises:

- involve a complex domain that can be subject to empirical investigation
- involve uncovering the relevant components of the domain, and the causal relationships between those components
- apply the use of models to represent and investigate phenomena
- are essentially knowledge work, usually involving multiple actors working in collaboration on a variety of interconnected tasks

In other words, the types of problems facing scientists and organizational managers are not all that different, the context of their investigations is similar, and the techniques they apply have much in common. Given these strong similarities it thus seems entirely appropriate to apply insights from the philosophy of science to the challenge of managing knowledge work in complex organizational environments.

2 Theory/Issues - The Disunity of Science and the 'Local View'

Historically, the ideal of the unity of science has been seen as the basis upon which the success of science has been built (Oppenheim & Putnam, 1958; Hempel, 1965). According to the unity of science view, scientific explanation gains its power from our discovery of general laws of nature – exceptionless rules that govern the way the world works. The task of science then is to uncover these laws and to present a single, unified picture of the world. By use of these laws we can explain and predict events, manipulate and control nature, and discover novel facts about the natural world. These laws describe the fundamental workings of the world, at the level of the basic building blocks of the constituents of the world, and are applicable at any time in any place in the universe. According to this ideal, science is seen as the quest for the ultimate rules according to which the universe runs, with everything following from these simple and elegant equations. As such this is a *top-down* view, since on this view all scientific explanation and prediction is based on deduction from universal, overarching laws.

The disunified view turns this standard view on its head – quite literally. The disunified approach denies that we must develop theories that describe the workings of the world from the top-down. Instead of emphasising overarching laws of nature, the disunified view builds our understanding of the world from the bottom-up, by identifying the relevant properties and processes at work in the world, and exploring them through the use of experiment and modelling.

This disunified approach has emerged from recent work in the philosophy of science, which has moved away from the approach taken by philosophers of science such as Popper (1959) and Kuhn (1970), who engaged in the quest to develop a general, overarching and unified account of what science is. Instead, recent work has looked more closely at the fine detail of science as it is practised. These fine details concern the complex methods by which scientific theories are developed, in terms of how scientists work, reason, experiment, collaborate, and so on. Thus Cartwright (1989) emphasises the importance of *causal capacities* in science, and Dupré (1993) explores the metaphysical implications of the disunity of perspectives that coexist across the range of sciences. The detailed work of Galison (1996, 1997) looks at the role of social dynamics and politics in the theoretical life of nuclear physicists. Hacking (1999) also explores these issues in some detail, showing how the social construction of the world does not entail losing contact with traditional epistemological ideals such as accuracy and truth. Finally, Kitcher (1993) develops a complex model of scientific reasoning in a collaborative environment, a model that factors in the interactions between different researchers in building up a detailed picture of knowledge production in group context.

As a result of these inquiries, a richer, more accurate, and more powerful conception of science has been developed, which incorporates the diverse methodologies of different scientific disciplines, as well as the divergent metaphysics implicit in these different modes of inquiry. This approach has revealed that the standard view is guilty of what some have termed *fundamentalism*, which can been defined as the "tendency to think that all facts must belong to one grand scheme" (Cartwright 1994: 221) or "the dogma that order, either discovered in or imposed on the world, is a fundamental condition of the possibility of knowledge about the world" (Clarke 1998: 2-3). According to Cartwright, we should resist the lure of fundamentalism and the associated ideals of beauty, unity, simplicity and universality, not simply because they assume a false picture of the world, but because they entail a poor methodological approach to science:

> The problem is that our beliefs about the structure of the world go hand-in-hand with the methodologies we adopt to study it. The worry is not so much that we will adopt wrong images with which to represent the world, but rather that we will choose wrong tools with which to change it. We yearn for a better, cleaner, more orderly world than the one that, to all appearances, we inhabit. But it will not do to base our methods on our wishes. We had better choose the most probable option and wherever possible hedge our bets. (Cartwright 1999: 12-13)

In contrast to the fundamentalist view, the thesis of the disunity of the sciences views the world of scientific theorising as consisting of many varied and distinct areas of research, each with different aims, different standards of success, different terminology and different methodologies. According to this disunified view, there may be nothing that unites the diverse range of sciences: no overarching laws, no grand unified theory, and no ultimate reduction to physics. Because of this we must adopt what can be termed the 'local view' of theories, and a corresponding 'localised methodology' of science. According to the 'local view' of scientific theories, theories are not necessarily globally applicable, but instead provide locally applicable models based on relevant causal and stochastic processes. A 'localised methodology' is a scientific methodology that is concerned with understanding and revealing the entities and causal relations in a particular situation, with a particular problem at hand, and thus particular explanatory or predictive aims. Instead of appealing to generalisations or laws for explanation or prediction, a localised methodology is based on the uncovering and modelling of local causal and stochastic processes, involving locally specified capacities, properties and kinds. Thus, according to the local view, theories are based on locally specific models derived from causal and stochastic processes, they involve case study oriented approaches (Yin, 1994), and they are 'bottom-up' structures – explanations based on local causal and stochastic facts.

In science a localised methodology can deliver powerful predictions and detailed explanations, by rejecting the central importance of fundamental laws, by being open to the possibility of disunity, and by focussing on solving particular problems in particular contexts rather than developing generally applicable theories. Such a methodology is particularly applicable in the realms that involve complex physical systems in complex environments. In these complex sciences, this methodological approach has significant heuristic power, which is derived from uncovering and modelling the properties and processes that underlie the complex systems (Aarons, 2001). This approach thus provides powerful insights for guiding our understanding, manipulation, and management of these systems.

For these reasons the localised approach seems particularly suitable for analysing other types of complex systems, such as those involved in complex organizational environments. Such an approach should be directly applicable to these settings, providing a means to develop a methodology for managing and supporting knowledge work.

The first crucial point to note here concerns the importance of resisting fundamentalism. The lesson is that just as one should avoid the lure of fundamentalism in good science, one should avoid fundamentalism in organizational management. That is, when analysing and managing complex organizational systems one should not assume that the facts of one particular case can be generalised to other cases. Similarly one should never assume that there is a grand overarching approach suitable for all circumstances. What this means is that there is no magic bullet for organizational management – there is no single approach to the design, development, and implementation of new systems that will apply to all organizational contexts.

Secondly, the disunified, localised approach entails that an understanding of an organizational environment should be generated from the bottom-up rather than the top-down. That is, any investigation into the workings of a system should be made by a thorough analysis of the components of the system, looking at how those components piece together, and the processes that flow between those components. For a knowledge management project the starting point of such an analysis would be to determine precisely what aspects of knowledge are relevant to that particular case, and to give an account of the factors underlying these knowledge components. This would involve assessing the relevant cognitive, social and pragmatic factors involved in that particular KM project. From such an analysis a complex picture can be developed, piecing together the details until they form a complex model of the system. This can then form the basis for a detailed set of models that represent the system accurately. It is this approach that forms the basis of the Bureau of Meteorology case study, discussed later in the paper.

This bottom-up approach is in contrast to one which builds understanding from the top-down, by trying to fit all aspects of the system to a single scheme. The advantage of the bottom-up method is that it does not assume there is a single way of modelling the system, and does not assume that the system can be modelled in its entirety by a single model. Thus a complex organizational system could be modelled by a set of distinct models, all representing different though perhaps not unrelated aspects of the broader system. As a result the set of models is not a nice neat arrangement that merges together to form a large unified picture. Instead, the picture is more of a dappled patchwork (Cartwright, 1999), which if looked at all at once seems a blurred mess. One can instead focus on small sections of the patchwork, seeing how the edges of the sections are badly frayed and clearly do not fit in with adjoining areas. Yet within each patch, in the local context it focuses on, there is a clear view of every detail required for all practical purposes.

One further advantage of this approach is that it can support a theory of collaborative knowledge work within a *realist* and *pluralist* metaphysical framework (as outlined in Cartwright, 1999). That is, this approach maintains a connection with real-world processes and properties, and results in models that represent these real processes and properties. Importantly, this framework acknowledges the significant social dimension of knowledge work in such organizational settings, while retaining the idea that social processes are deeply connected to real properties and processes. Thus, in a more general way, applying a disunified methodology will make it possible to build a theoretical framework for supporting knowledge work that is grounded in reality, but also incorporates the relevant social, practical, and pragmatic concerns that are central to the fundamental tasks of organizational environments. Such a project is clearly beyond the scope of this paper. However by the end of the paper it should be clear that the outlined methodology is suitable for achieving this goal.

In many ways this approach should be seen as a development of the task-based knowledge management framework outlined by Burstein and Linger (2003). Indeed, the conception of knowledge work used in this paper has been strongly informed by the diverse range of KM projects and field studies discussed by those authors. According to the task-based framework, in a particular organizational setting the relevant knowledge is always situated in a specific context, so an organization benefits from a knowledge management system when such a system is based on knowledge workers' understandings of their specific tasks. The upshot of this is that in order to achieve a tasked-based understanding of knowledge work one must undertake a detailed study of the nature of the task. This in turn appeals to the situated activity theory of Iivari and Linger (2000), according to which knowledge work is built around an *activity system.* What the disunified view adds to these approaches is an explicit methodology for uncovering the particular complexities and relevant properties and processes within an organizational context.

3 Methods / Procedures / Results – Application: Forecast Streamlining and Enhancement at the Bureau of Meteorology

To illustrate the significance and power of the disunified approach, the case study of the Australian Bureau of Meteorology (BoM) is particularly illuminating. In this case the disunified approach has allowed for a detailed analysis of the nature of work conducted by meteorological forecasters, and as a result is contributing to the ongoing development and implementation of new and improved forecasting systems. This work builds on the knowledge management framework developed by Linger et al (2000), and has led to a rich understanding of the meteorological forecasting process as knowledge work.

Fundamental to this approach is recognition of the unique characteristics that make up the varied tasks involved in preparing a weather forecast. In particular, this includes the specific contextual details within which the work activity is conducted:

> The work of a meteorological forecaster is complex, characterised by uncertainty, incomplete information, multiple sources and a great variety of data, and strict timelines all overlaid by a legal regime. Forecasters are required to exercise judgement because science is often inadequate at the level of detail required by specific forecasts. In such an environment, work activity assumes not only task performance (constructing a forecast) but also the review and re-

> assessment of the work done in order to understand and learn from the experience. Moreover, meteorological forecasting is a continuous and collaborative process that is geographically distributed and involves a number of meteorologists within a forecast cycle and between cycles. Such work is unambiguously knowledge work. (Linger et al, 2000: 123)

The way forecasts are constructed by forecasters has remained largely unchanged since the initial revolution in weather forecasting, made possible by profound improvements in accurate satellite imaging and powerful numerical modelling from the 1950s onwards (Fishman & Kalish, 1994). Although new tools have been built over recent years to incorporate new inputs (such as an integrated data viewer and forecast preparation package in the Australian system), these tools are essentially designed to facilitate the existing ways of forecasting rather than to improve the forecasting process. As such, these tools are limited in their scope: they cannot meet the new range of challenges for meteorology that have arisen as a result of recent scientific and technological developments, as well as a renewed public (and legal) interest in accurate and informative weather forecasting.[1]

It is in this context that the conception of forecasting as knowledge work has facilitated the development of a knowledge management framework known as the "Mandala Project" (Linger et al, 2000; Linger et al, 2001). This framework has provided a conceptual model of an improved forecasting system which incorporates many recent scientific and technological developments that present significant new challenges for meteorology. These include

- Adapting and incorporating newly available Numerical Weather Prediction (NWP) model outputs into a diverse range of general and specialised weather products, which need to be prepared within a fixed time frame and conform to a particular fixed format;
- Managing the volume of diverse material that needs to be referenced in order to prepare the forecasts;
- Managing the data holding to make relevant forecast aids, guidance and alerts available during the forecast process;
- Improving the accuracy of forecasts, and the justification of those forecasts;

[1] Much of the impetus for improving forecasting systems has arisen from two recent investigations into high profile cases of perceived forecasting problems: the 1998 Sydney to Hobart yacht race tragedy (Bureau of Meteorology, 1999a) and the 1999 Sydney hailstorm investigation (Bureau of Meteorology, 1999b).

- Improving decision support to assist meteorologists to exercise judgement and to use their experience and tacit knowledge to overcome the limitations of scientific knowledge. (Linger et al 2001)

The knowledge management architecture that emerges from the Mandala Project is represented by the "pill" diagram in figure 1.

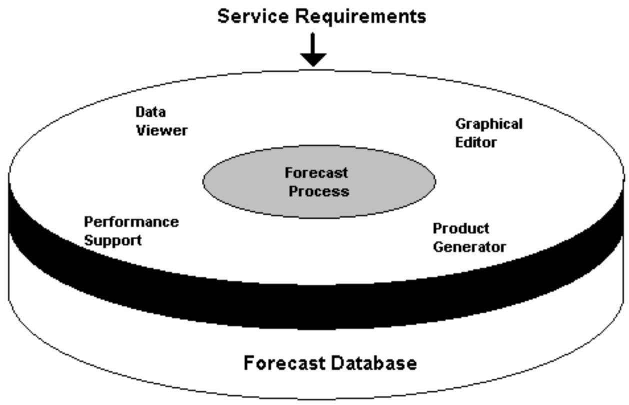

Fig.1. Mandala Project model for future forecasting systems. Adapted from Linger et al (2001)

Adopting the recommendations that emerged from the Mandala Project, the BoM has undertaken what it calls the "Forecast Streamlining and Enhancement Project" (FSEP). FSEP essentially involves implementing the architecture represented in figure 1, through the development of new forecasting systems and the creation of a centralised forecast database. More significantly (and more problematically), it is clear that the adoption of this new approach involves a fairly significant shift in work practices. This shift entails a number of significant changes to the way forecasting work is conducted (Bell, 2003).

The most significant change is that the work of the forecaster will change from being *product centred* to being *information centred.* What this means is that rather than the forecaster working on specific types of forecast *products*, such as marine forecasts, aviation forecasts or particular region forecasts, the forecaster will instead work on predicting the values of particular weather *elements* (parameters) for a given region. The FSEP ideal is that once the various elements have been forecast and input into a central forecast data base (FDB), the individual products can then be generated automatically from the FDB. Such a system will incorporate a number of features which will dramatically improve the accuracy and efficiency of the forecast process. These features include: minimal checking of products; automatic update of products; auto text generation to plain language; and intelligent alerting and decision support. A comparison of the FSEP approach with the present approach is presented in Table 1 below.

Table 1. Comparison of forecast systems.

Present Forecasting System	FSEP System
Product centred.	Information centred.
Forecasting work organized around the preparation of forecast products.	Forecasting work organized around the prediction of forecast elements.
Multiple forecast products, generated and edited *individually* and *manually* by the forecaster.	Multiple forecast products, all *automatically* generated from the forecast data base.
Time consuming manual preparation and checking of products by forecaster.	Auto text generation with minimal product checking by forecaster.
Manual updating of products.	Automatic updating of products.
A fragmented approach, with seven different forecast divisions, and a number of inconsistencies in terminology and methodology between the divisions.	An integrated, nationally consistent approach.
No centralised forecast database. Meteorological data is largely distributed across the organization, and mostly held in forecasters' memory.	A centralised forecast database (FDB), holding meteorological data for the whole organization.
Data viewing and editing facilitated by separate systems, with minimal record keeping.	Integrated data viewing and editing, feeding directly into the FDB, along with relevant metadata.
Most products are text based.	Increasing number of graphical based products.

To this point the Mandala Project (Linger et al 2000) has provided an excellent overall specification for the project. It gives a clear indication of the particular objectives of the FSEP system, how they relate and integrate with work practices, and sets up the basis for design criteria for this system. As such it presents a good model of the endpoint of the transition process. However the weakness of this approach is that this model does not help sufficiently with the precise details of system development and implementation, and has little to say about how to manage the transition process from the existing system to the streamlined approach. The model as it is presented in Figure 1 is extremely limited in its scope, since it really only represents the key IT components of the new forecasting system – it does not situate the human forecaster in the model, and it does not explain how the different pieces of the model actually fit together. There is thus a need to supplement the existing model, in order to provide guidance for the future implementation of the FSEP strategy, and to incorporate these factors into the broader FSEP vision.

This is where the disunified, bottom-up approach proves its strength, as better alternative to the strategy adopted up to this point. In contrast to the approach of the Mandala Project, which essentially adopts a top-down analysis to build a globally applicable model, the disunified approach starts at the bottom, with the particular details of the specific tasks, product types, actors and elements. This approach suggests that a detailed conception of the FSEP system be built up from these components, and that development and implementation also proceed in this way from the bottom-up.

It is important to note here that this approach is not incompatible with the top-down approach. In fact, both approaches are necessary, as the top-down approach provides the scope and broad specifications within which the particular analysis and implementation given by the bottom-up approach will occur.

As it happens, what is now occurring at the BoM with the implementation of FSEP is very much along the lines of the approach just described: a number of modules are being developed that conform more or less to the specifications of the FSEP vision. These modules include the Tropical Cyclone forecasting system (TC), the Thunderstorm Interactive Forecast System (TIFS), and the Australasian Marine Forecasting System (AMFS). These modules are in the process of being introduced into a number of the divisions of the BoM, with each implementation involving the integration of the particular module into the present forecasting system, and the customization of that module to meet particular regional requirements. The essential problem with this is that each region constructs forecasts in a slightly different way, due to differing standards, requirements, and practices between each region (Bally et al, 2004). The modules themselves are also quite separate entities, which are used separately to perform quite distinct tasks. Even though they are integrated into the forecasting process, they do not merge together to form a unified system. In fact, the work of the forecaster is largely split up into numerous, individual tasks, each performed separately, using different components of the system to perform each task. As such, the forecasting system use by any particular forecaster is really a conglomeration of numerous modules and subsystems, working quite separately. All the systems really have in common is that they are used to construct forecasts, which are then entered into the central forecast publishing system, the Australian Integrated Forecast System AIFS (Kelly & Gigliotti, 1997). In effect, this means that the forecasting modules are essentially disunified – they are separate, non-integrated entities. Thus the knowledge work that characterises forecasting is also disunified, since it is comprised of numerous disparate tasks, requiring different techniques and abilities, and utilising different terminologies and systems.

It is for this reason that the *only* effective methodology for both analysis and implementation is a bottom-up approach. This is because only in this way can the particular differences between the different forecasting tasks successfully be accounted for. Trying to conceive of forecasting as a single unified activity would make one guilty of fundamentalism, and would lead one to a false and unworkable picture of the forecasting process.

There is, however, some unity to the process of developing the FSEP modules. For although the modules are being implemented and customised differently in each region, the development of each module is occurring centrally and consistently. This is being achieved by ensuring that the modules are developed according to a specified set of standards, including the use of Java within a particular standard environment called VisAD (http://www.ssec.wisc.edu/~billh/visad.html [February 12, 2004]), the use of consistent interfaces and GUI standards, and the use of standard data models, all in an environment built to allow open and consistent data sharing. Most importantly, and most relevantly as far as the disunified view is concerned, the FSEP working group within the Bureau have explicitly adopted the methodology of *agile development*, also known as *eXtreme Programming* or XP (Beck, 2003), for the development of these FSEP modules. The methodology of XP embraces the following set of principles:

- Whole Team participation
- Customer Tests
- Small Releases
- Simple Design
- Pair Programming
- Test-Driven Development
- Design Improvement
- Continuous Integration
- Collective Code Ownership
- Coding Standard

(http://www.xprogramming.com/xpmag/whatisxp.htm [February 21, 2004])

What is so significant about this approach, for the purposes of this paper, is that it explicitly adopts the policy of incremental change in systems development, continuously building up a working system to incorporate the immediate needs of the system users. In this case the users are the meteorologists, who work together with the programmers by giving rapid and constant feedback on each software release. In this way each module is built incrementally from the bottom-up.

At present this approach has resulted in each module being a separate system, developed fairly independently from each other (though often by the same team of developers), and implemented independently and in different ways across the seven different forecast divisions. The main worry with this approach is that the XP methodology will not deliver the promise of developing a large, integrated, and fully working system – a concern voiced by a number of critics of XP (eg Stephens, 2003; Boehm & Turner, 2003). The problem is that this approach may result in a large number of distinct and potentially incompatible systems, unless great care is taken to ensure that the strict architecture is adhered to at all stages of development.

A further, and perhaps more serious concern, is that the XP approach is just not suitable for one of the main elements of the FSEP system, the integrated Forecast database (FDB). In particular, it is not clear that the FDB can be developed in a piecemeal way, because it must be designed carefully up front in a way that explicitly goes against the XP mantra of no Big Design Up Front (BDUF) (http://xp.c2.com/BigDesignUpFront.html [February 12, 2004]). This problem has been borne out in practice, as very little progress has been made on the development of the FDB, while development and deployment of the modules has been proceeding at a constant rate.

In order to investigate the FSEP development and implementation process more thoroughly, and in accordance with the disunified approach, a detailed study of one of the FSEP modules has begun. This study is looking at the ongoing adoption of the Australasian Marine Forecasting System (AMFS) (Kelly et al, 2003) in a number of Bureau regions. The AMFS module incorporates many of the FSEP principles in its design: it is a multi-element system; it builds on Numerical Weather Prediction (NWP) guidance with forecaster 'value adding'; it incorporates a graphical editor and utilises a central database; and it includes auto product generation. To date, AMFS has been adopted to different degrees in a number of regional offices. AMFS has been operational in Tasmania since December 2002, with the coastal waters forecast currently available on the web with graphical output. South Australia uses AMFS to provide weekend inland waters forecasts with graphical output. New South Wales has recently begun using the module to provide weekend boating forecasts for the Sydney region. There are also plans to introduce the module into the Victorian regional office, where it currently is in the setup planning phase.

The idea of AMFS is to produce multiple Marine Forecasts, tailored to regional requirements, using a single consistent process. Like the other FSEP modules, AMFS is a single application that can be tailored to regional requirements (through the use of flexible XML, allowing for customizable colours, graphics presentation and automatic text generation.). AMFS works by feeding inputs from guidance data (NWP data, satellite images, and direct observation data) directly into an integrated graphical editor. The forecaster can then choose to manually adjust the representation in the data editor if need be. This could occur, for example, if there is a known bias in the results of the NWP inputs, such as a tendency to bias wind predictions towards southerly winds. This stage allows the individual forecaster to ‘add value’ to the forecast, incorporating their own expert knowledge into the forecast process. Once the forecast has been finalised by the forecaster all it takes is a push of a button to save the results in the FDB, and to automatically output a diverse range of forecast products in many different forms of media. Figure 2 is an example of the automatic web output from AMFS in Tasmania.

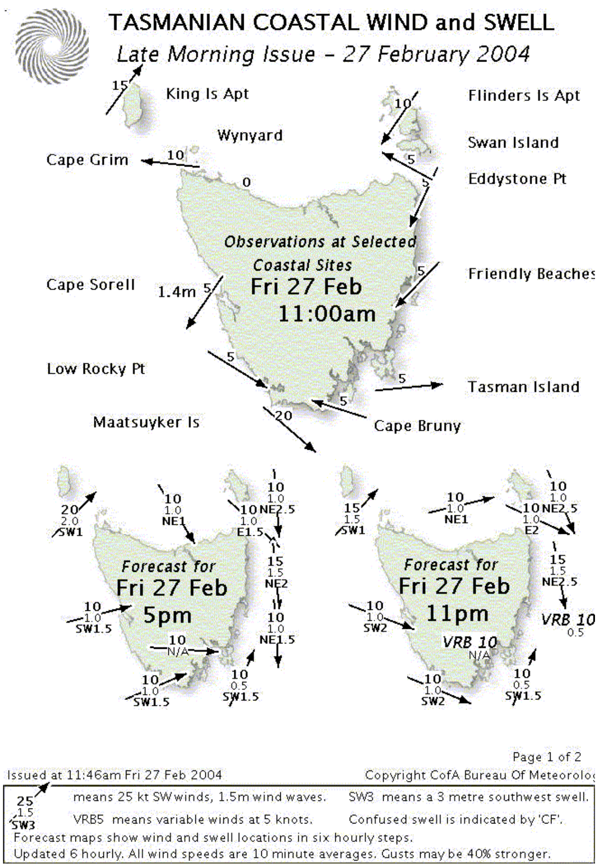

Fig. 2. Sample output from AMFS

Although the detailed study of the AMFS module is still in its early stages, at present it is clear that the adoption of AMFS has been at least partially successful (C. Ryan & J. Kelly – personal communication). The evidence for this success comes both from the personal accounts of developers and forecasters interviewed as part of this study, and from the public popularity of the web-based products generated by the AMFS module. To follow this up, a precise ethnographic study and associated data collection will be conducted, involving observational studies of forecasters at work, looking in detail at the adoption and integration of AMFS into the forecasters’ daily work practices. This study will delve into the messy details of the tasks performed by forecasters, along the lines of Schultze's (2000) ethnographic investigation into knowledge work.

Central to the approach taken here is the assumption that the best way to investigate forecasting as knowledge work is to build an understanding from the bottom-up, thereby building a model of the forecasting process based on the tasks performed by the forecasters, incorporating the particular peculiarities unique to each forecasting context. It is only through an investigation such as this that the exact design specifications and precise implementation requirements can be determined, since each particular case is in effect unique.

4 Conclusions

This paper began with an outline of the disunity of science view, and argued that the methodological insights that come from taking a disunified view of science can be applied directly to the domain of knowledge management and the task of managing knowledge work. The case study of the FSEP project at the Australian Bureau of Meteorology has demonstrated that these insights can be applied to complex organizational environments. In particular, this case has demonstrated two significant points:

- the limitations of adopting a purely top-down approach to develop a knowledge management framework in a particular knowledge work context
- the strength of adopting a bottom-up, disunified approach to analysing an organizational environment, especially for the development and implementation of systems for supporting knowledge work

Although further work is needed to support these conclusions with empirical rigour, there is a strong *prima facie* conceptual case in favour of adopting the suggested approach. In particular, this project has so far demonstrated that developing an explicit account of the processes and properties central to an analysis of knowledge work can provide much needed support for the practical tasks of knowledge management.

One further conclusion that arises from this approach is that, although the disunified view entails that there can be no general fix-all solution for supporting knowledge work in all contexts, the *methodology* of the disunified approach clearly is generalizable. As such it goes at least part of the way to meeting Drucker's (1999) challenge of improving knowledge worker productivity.

5 References

Aarons, J. (2001). *Thinking Locally: A Disunified Methodology of Science.* [Unpublished PhD manuscript, School of Philosophy and Bioethics, Monash University]

Bally, J., Boneh, T., Nicholson, A., and Korb, K. (2004) Developing An Ontology for the Meteorological Forecasting Process. [Unpublished manuscript – submitted to the 2004 IFIP International Conference on Decision Support Systems (DSS2004), Prato, Italy, 1st-3rd July 2004.

Boehm, B. & Turner, R. (2003). *Balancing Agility and Discipline: A Guide for the Perplexed.* Boston: Addison-Wesley

Bureau of Meteorology (1999a). Preliminary report on meteorological aspects of the 1998 Sydney to Hobart yacht race. http://www.bom.gov.au/inside/services_policy/marine/sydney_hobart/prelrept.html [February 20, 2004]

Bureau of Meteorology (1999b). Report by the Director of Meteorology on the Bureau of Meteorology's Forecasting and Warning Performance for the Sydney Hailstorm of 14 April 1999. http://www.bom.gov.au/inside/services_policy/storms/sydney_hail/hail_report.shtml [February 20, 2004]

Beck, K. (2003). *Extreme Programming Explained: Embrace Change.* Boston: Addison-Wesley

Bell, I. (2003) FSEP Principles. [Bureau of Meteorology – internal technical report]

Burstein, F. & H. Linger (2003). Supporting Post-Fordist Work Practices: A Knowledge Management Framework for Supporting Knowledge Work. *Information Technology & People* Special Issue on Organizational Implications of Knowledge Management Systems.

Cartwright, N. (1989). *Nature's Capacities and their Measurement*. Oxford: Clarendon Press.

Cartwright, N. (1999). *The Dappled World.* Chicago: University of Chicago Press.

Clarke, S. (1998). *Metaphysics and the Disunity of Scientific Knowledge.* Aldershot: Ashgate.

Cortada, J. W. (ed.) (1998). *Rise of the Knowledge Worker.* Boston: Butterworth-Heinemann.

Davenport, T. & L. Prusak (1998). *Working Knowledge: How organisations manage what they know*. Harvard Business School Press.

Drucker, P. F. (1999). Knowledge Worker Productivity: The Biggest Challenge. *California Management Review,* Vol.1 No. 2, pp. 79-94.

Dupré, J. (1993). *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Cambridge, Mass: Harvard University Press.

Fishman, J. & Kalish, R. (1994). *The Weather Revolution: Innovations and Imminent Breakthroughs in Accurate Forecasting.* New York: Plenum Press.

Galison, P. & D. Stump (eds.) (1996). *The Disunity of Science: Boundaries, Contexts, and Power*. Stanford: Stanford University Press.

Galison, P. (1997). *Image and Logic*. Chicago: University of Chicago Press.

Hacking, I. (1999). *The Social Construction of What?* Cambridge, Mass.: Harvard University Press.

Hempel, C. G. (1965). *Aspects of Scientific Explanation*. New York: The Free Press.

Iivari, J. & H. Linger (1999). Knowledge work as collaborative work: A situated activity theory view. *Proceedings of the Hawaiian International Conference on Systems Science* (HICSS'32).

Iivari, J. & H. Linger (2000). Characterizing Knowledge Work: A Theoretical Perspective. *Proceedings of the Americas Conference on Information Systems AMCIS'2000.*

Kelly, J. & Gigliotti, P. (1997), 'The Australian Integrated Forecast System (AIFS): Overview and current status', Preprints, 13th Int. Conf. on IIPS for Meteorology, Oceanography and Hydrology, pp. 141-144.

Kelly, J., Bally, J., Ryan, C., and Thomas, D. (2003). Australian Marine Forecasting System. Preprints, Conf. Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology. Long Beach, Amer. Meteor. Soc., http://ams.confex.com/ams/annual2003/19IIPS/abstracts/55274.htm [July 2003]

Kitcher, P. (1993). *The Advancement of Science.* Oxford: Oxford University Press.

Kuhn, T. (1970). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

Linger, H., Burstein, F., Kelly, J., Ryan, C., and Gigliotti, P. (2000). Creating a learning community through knowledge management: the Mandala Project. Proceedings of the IFIP WG 8.3 conference “Decision support through knowledge management”, 9-11 July 2000, Sweden, pp. 122-36.

Linger, H., Burstein, F., Ryan, C., and Kelly, J. (2001). Implementing a Knowledge Management System: the Case of Meteorological Forecasting. In Burstein and Linger (eds.) (2001) *Knowledge Management for Information Communities,* pp. 139-153. Melbourne: Australian Scholarly Publishing,

Oppenheim, P. and Putnam H. (1958). Unity of Science as a Working Hypothesis. In Feigl, H., Scriven M., and Maxwell G. (eds.) *Minnesota Studies in the Philosophy of Science. Volume II.* pp. 3-36. Minneapolis, MN: University of Minnesota Press.

Popper, K. (1959). *The Logic of Scientific Discovery.* London: Hutchinson.

Ruth, D. P. (2000). Models, forecasters, and interactive forecast preparation in the new millennium. *Preprints 16th Int. Conf. on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology*, pp.82-85. Long Beach: Amer. Meteor. Soc.

Ruth, D. P. (2002). Interactive Forecast Preparation – the future has come. *Preprints 18th Int. Conf. on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology*, pp. 20-22. Orlando: Amer. Meteor. Soc.

Ryan, C. (2003). The Forecast Streamlining and Enhancement Project. [Bureau of Meteorology – internal technical report]

Schultze, U. (2000). A Confessional Account of an Ethnography about Knowledge Work, *MIS Quarterly,* Vol 24, no 1, pp.1-39.

Stephens, M. & Rosenberg, D. (2003). *Extreme Programming Refactored: the Case Against XP*. Berkeley, Calif.: Apress; Berlin: Springer.

Yin, R. K. (1994). *Case Study Research: Design and Methods*. (Second Edition) Thousand Oaks, California: Sage.