1. Theory Building in the Information Systems Discipline: Some critical reflections

Ron Weber
Monash University

Abstract

This chapter articulates criteria for evaluating the quality of a theory. It also shows how the criteria can be used to pinpoint the strengths and weaknesses of a theory that, based upon citation evidence, appears to have had a significant impact on other researchers within the information systems discipline. Aside from their evaluation purposes, the criteria are also intended to inform researchers who are seeking to build high-quality theory to account for some type of phenomena within the information systems domain.

> Nothing is as practical as a good theory. (Lewin, 1945, p. 129)

Introduction

For many researchers, the development of theory within their discipline is the central goal—the 'jewel in the crown'—of their research endeavours (for example, Eisenhardt, 1989). By articulating high-quality theory, they believe they are more likely to enhance their own knowledge of, other scholars' knowledge of and practitioners' capabilities to operate effectively and efficiently in their domain of interest.

In spite of the importance ascribed to theory by many researchers, the development of theory has been a relatively neglected feature of research within the information systems (IS) discipline. As a result, in the late 1990s and early 2000s, several editors of major journals appealed for more theoretical contributions to be made to the discipline (for example, Zmud, 1998). Moreover, only recently have scholars within the discipline begun to focus on how *high-quality* theory can be developed (for example, Gregor, 2006). Nonetheless, in many respects, the process that should be used to develop high-quality theory remains an arcane affair.

In this chapter, I propose a set of normative criteria that can be used to evaluate the quality of a theory. I show how these criteria can be used to pinpoint the strengths and weaknesses of a theory that has been proposed within the information systems discipline—one that according to citation evidence has had a significant impact on many researchers within the discipline. In this way, I seek to illustrate the usefulness of the criteria. As a related matter, I seek to show also that the criteria provide useful guidelines for *developing* theory.

The structure of the chapter is as follows. First, I briefly define some ontological constructs that enable me to define the meaning I ascribe to the term 'theory' and to articulate the criteria I propose for evaluating the quality of a theory more precisely. Next, I explain the meaning I ascribe to the term 'theory'. I then describe the criteria that I propose for evaluating the quality of a theory. Subsequently, I attempt to show the usefulness of these criteria by applying them to the evaluation of an important, extant information systems theory. Finally, I provide some brief reflections and conclusions.

Some Basic Ontology

To provide the basis for my analyses in the sections that follow, this section provides a brief (and somewhat informal) explanation of some fundamental ontological constructs. These constructs are derived from a formal ontological theory proposed by Bunge (1977, 1979).

- *Thing*: The world is made of things. Things can be *substantial* or *concrete* (for example, people or buildings); alternatively, they can be *conceptual* (for example, sets and functions). In this chapter, my focus is primarily on *concrete* things.
- *Composite thing*: Some things are made up of other things (for example, a team, which is a composite thing, is made of team members, which are its components).
- *Property*: All concrete things in the world possess properties (there are no formless things). Similarly, all properties in the world attach to some thing (properties do not exist in isolation from things). Properties are *not* things; they are separate ontological constructs that describe different elements in the world.
- *Class*: Things that possess at least one property in common constitute a class of things.
- *Attributes*: We 'know' about properties of things in the world through our *perceptions* of them. These perceptions might be more or less true. The way

in which we perceive a property at a point in time (our representation of it) is called an attribute.

- *Types of attributes*: Various types of attributes exist
  - *intrinsic attributes* represent properties of individuals (for example, the height of a person)
  - *mutual attributes* represent properties of two or more things (for example, the date on which one person was married to another person)
  - *emergent attributes* are attributes of composite things that do not belong to their components but nonetheless are related in some way to attributes of their components (for example, the work productivity of a team has no meaning in terms of each team member, but it is related in some way to the productivity of each team member)
  - *attribute in general*: attributes in general belong to a class of things (for example, all humans possess the attribute called 'height').
  - *attribute in particular*: particular things in a class of things possess attributes that have particular values (for example, the thing called 'John' in the class of things called 'people' possesses the particular attribute 'height is 180 centimetres').
- *State*: A vector of attributes in particular represents a state of a thing (its attributes in general along with their associated values). Some states are deemed *lawful* (they obey natural or human-made laws); others are deemed unlawful.
- *Event*: An event that a thing undergoes is represented by a change from one of its states to another of its states. Some events are deemed *lawful*; others are deemed unlawful. If an event has an unlawful beginning or end state, the event will be unlawful. Some events are unlawful, however, even when their beginning and end states are *lawful*. For instance, 'alive' and 'dead' are lawful states of a human thing. The event represented by the state change from 'alive' to 'dead' is lawful; the event represented by the state change from 'dead' to 'alive' is unlawful.
- *Interaction*: Two things interact when the history of one thing, which is manifested as changes in its states (or attributes in particular), is not independent of the history of the other thing.

In the sections below, I use these constructs to explain the nature of theory and criteria that can be used to evaluate the quality of a theory.

Nature of Theory

Different researchers often ascribe different meanings to the term 'theory'. For instance, Gregor (2006) examines five ways in which the term 'theory' has been used in the literature. I do not agree that the term 'theory' covers all the meanings she canvasses (in my view, those she calls Type-I and Type-V 'theories' are not theories). For this reason, in this section I explain the meaning I ascribe to the term 'theory'.

By *theory*, I mean a particular kind of *model* that is intended to *account* for some subset of *phenomena* in the world. A theory is a *social construction*. It is an artefact built by humans to achieve some purpose. It is a *conceptual* thing rather than a concrete thing.

By *phenomena*, I mean changes in the *attributes in particular* of some things in the world. The subset of phenomena in the world that the theory is intended to cover is called the *focal phenomena*.

By *account*, I mean a theory assists its users to *predict* and/or *explain* its focal phenomena. Some researchers ascribe another purpose to theories—namely, to facilitate human *understanding* of the theory's focal phenomena. I do not see how *explanation* of focal phenomena can occur, however, without first *understanding* the focal phenomena. For this reason, I intend the purpose of explanation to encompass the purpose of understanding.

By *model*, I mean a *representation* of something else (phenomena) in the world. Theories are, however, particular kinds of models (see section four below). All theories are models, but not all models are theories. A model must satisfy particular conditions before I deem it to be a theory (see below).

Framework for Theory Evaluation

In this section, I argue that a theory must be evaluated from two perspectives. The first is the 'parts'—the evaluation must focus on the quality of the individual components that make up the theory. The second is the 'whole'—the evaluation must focus on the quality of the theory considered *in toto*. Both forms of evaluation are important in assessing the quality of a theory. Clearly, it is unlikely that the quality of the whole will be high if the quality of the parts is not high. Nonetheless, high-quality parts are a *necessary* but not *sufficient* condition for a high-quality whole. To the extent a *model* satisfies these criteria, it can be deemed a *theory*.

Parts

A theory has three parts (or components): its constructs, its associations and its boundary. When evaluating a theory, the focus initially should be on the quality of these parts. The following subsections explain the nature of each part. They also describe criteria that can be used to evaluate how well a researcher has articulated each of the parts.

Constructs

A *construct* in a theory represents an *attribute in general* of some *class of things* in the focal domain (as opposed to a *particular attribute* of a *specific* thing). The classes of things to which attributes in general pertain ought to be defined precisely to ensure the *meanings* of each class and the things in each class are clear. Otherwise, the meanings of the attributes in general that attach to the classes of things are unlikely to be clear. Attributes do not float in the ether; they always attach to things. As a first step in clarifying the meaning of an attribute, therefore, the thing to which it attaches needs to be made clear.

Once the meanings of the classes of things that a theory covers are clear, the nature of each attribute in general that pertains to a particular class ought to be defined precisely. Unless the meanings of the attributes in general are clear, the meanings of any associations among the attributes in general cannot be clear. Moreover, developing credible (valid and reliable) empirical indicators of the attributes in general will be difficult (if not impossible).

Associations

An association between two constructs in a theory shows that a *history of an instance* of at least one of the constructs is conditional on a *history of an instance* of the other construct. In other words, at least *one change in the value of an instance of one construct* is somehow related to at least *one change of value in an instance of the other construct*.

If two constructs represent different attributes in general of a *single* class of things, any association between them means the two attributes are lawfully related. In other words, for at least one instance of a thing in the class, a change in the value of at least one of the two attributes is related to a change in the value of the other attribute.

If two constructs represent different attributes in general of two *different* classes of things, any association between them means at least one instance of a thing in one class *interacts* with at least one instance of a thing in the other class. In other

words, the histories of the two things are not independent of each other. The nature of the interactions between the two things is manifested in the attributes that are related.

Associations can be specified with varying levels of precision (Dubin, 1978).

- Two constructs are simply shown to be related to each other, but neither the 'sign' nor the 'direction' of the association is shown. In other words, the association does not imply causality, nor does the association indicate whether a positive or negative change in the value of an instance of one of the constructs is associated with a positive or negative change in the value of an instance of the other construct.
- The sign of the association between two constructs is shown, which implies that changes in the values of an instance of one of the constructs are positively or negatively correlated with changes in the values of an instance of the other construct.
- The 'direction' of the association between two constructs is shown, which implies causality (some changes in value of an instance of construct A cause a change in value of an instance of construct B) or at least a time series of value changes (a change in the value of an instance of construct A precedes a change in the value of an instance of construct B).
- A functional association is shown between two constructs. In other words, the amount of change that occurs in the value of an instance of one construct as a result of a change that occurs in the value of an instance of another construct.

To the extent that the nature of the associations among constructs in a theory can be specified more precisely, empirical evaluations of the theory can be done more precisely. Moreover, to the extent that empirical tests of the theory support the existence of the associations, the theory has greater predictive and/or explanatory power.

When one of the associations in a theory shows directionality, the theory is a *process* theory (Markus and Robey, 1988). The reason is that the directionality means a change in the value of one of the constructs in the association *precedes* a change in the value of the other construct in the association.

Boundary

The boundary of a theory circumscribes the state space and event space of things for which the theory is posited to have predictive and/or explanatory power. A first step in specifying the boundary of a theory, therefore, is to be

clear about the class or classes of things that the theory covers and the attributes in general of this class or these classes of things that the theory covers. The constructs in the theory can then be defined precisely.

Unless the constructs in a theory are defined precisely, the boundary of a theory cannot be defined precisely. Often theories cover only certain values of each construct. If the constructs are not clear, the attributes of things in the class or classes of things covered by the theory will not be clear. As a result, the values of the attributes covered by the theory cannot be specified clearly.

Even if the constructs in a theory are defined precisely, however, researchers must still specify the *values* of each construct for which the theory is posited to hold. In the context of the theory, these values underpin the 'lawful states' of the theory—the states of things that the theory covers.

Similarly, researchers must specify the *events* that their theory covers. They must consider possible changes in the values of constructs (state changes) and whether the associations they posit exist among the constructs in their theory hold when such changes occur. In some cases, two states that a construct might assume might be covered by the theory, but a transition between these states might *not* be covered by the theory. Those state changes that are covered by the theory constitute the 'lawful events' of the theory.

When articulating the boundary of their theory, researchers must also consider whether their theory covers different *combinations* of construct states and different *combinations* of or *sequences* of construct events. The behaviour of constructs cannot be considered in isolation; instead, the implications for the theory of states of and events in subsets and sets of constructs must be considered.

The boundary of a theory must be specified precisely if precise empirical tests of the theory are to be undertaken. In the absence of a precise specification of a theory's boundary, researchers might test the theory unwittingly in an inappropriate context (one that falls outside its boundary).

Whole

A theory has *emergent* attributes—attributes of the theory *as a whole* rather than attributes of its parts. Many such attributes exist, and researchers often differ in their views of the significance they ascribe to each attribute. Nonetheless, some emergent attributes have widespread acceptance among researchers as being significant in the evaluation of the quality of a theory. The following subsections explain the nature of these attributes and describe criteria that can be used to evaluate the extent to which a theory possesses them.

Importance

The importance of a theory is often assessed via judgments made about the importance of its focal phenomena. Usually, there is little point to having a theory with rigorously specified constructs, associations and boundary conditions if it addresses uninteresting phenomena. The focal phenomena might be deemed important from the viewpoint of practice (improving the effectiveness and efficiency of some entity's activities). They might also be deemed important from the viewpoint of research. Potentially, enhanced understanding of the focal phenomena will provide key insights that will enable theoretical or empirical progress to be made on some problem within a discipline.

Ex ante, it might be difficult to judge the importance of a theory. At the outset, its potential impact on researchers and practitioners might be difficult to assess. Moreover, sometimes a theory provides insights that were not anticipated when it was first articulated. Such insights arise only when researchers engage with the theory and use it as the basis for their empirical work.

Ex post, however, various metrics are available to assess the importance of a theory. For example, the extent to which a theory is cited by researchers provides an indicator of its impact on their work and thus its likely importance to them. Similarly, whether a theory is cited in practitioner publications or underpins consulting work provides an indicator of its importance.

Citation evidence must be treated cautiously when it is used as a proxy for the importance of a theory. Some theories are appealing to researchers because they are relatively simple to test empirically. They are perceived as an easy route to journal publications. Whether they provide deep insights into the phenomena they cover, however, is another matter.

Novelty

The extent to which a theory is novel appears to be an important factor in determining: a) the value ascribed to it by researchers; and b) the likelihood that papers describing the theory will be accepted for publication in major journals (for example, Mone and McKinley, 1993). In short, judgments about a theory's novelty and judgments about its contributions to knowledge appear to be closely related. Moreover, the *importance* of a theory and its deemed novelty appear to be closely related. Nonetheless, some theories might be considered novel but cover phenomena that are regarded as unimportant.

Weber (2003) describes a number of ways in which a theory might make novel contributions to a discipline. For instance, a theory's focal phenomena might not have been covered by prior theories, or the focal phenomena might be framed or conceived in a different way. In essence, the theory being proposed in

these cases is a new theory. Alternatively, an existing theory might be modified by adding and/or deleting constructs and associations, defining constructs and associations more precisely or specifying the boundary of the theory more precisely.

A theory will also be deemed novel (perhaps after some time has elapsed) to the extent it changes the *paradigms* used by researchers to investigate phenomena within their discipline (Kuhn, 1996). It will command the attention of researchers if it provides a way of resolving 'anomalies' within their discipline—that is, empirical observations of phenomena that existing theories are unable to explain or predict. It will also command the attention of researchers if it enables them to 'see' or conceive of new and interesting phenomena (phenomena that previously escaped their attention) or re-conceptualise existing phenomena in new and interesting ways. Such theories break the cycle of 'normal science' within a discipline and set a new path for the discipline to follow.

The *quality of the rhetoric* used by researchers to describe their theories also appears to be an important factor in determining the extent to which their theories are deemed novel (Locke and Golden-Biddle, 1997). Because science is a social phenomenon, researchers have to convince their colleagues that their work has value. In this light, the arguments a researcher uses to expound their theory's novelty must be crafted carefully; otherwise, the theory's contribution to knowledge might be overlooked.

After analysing 82 papers published in the *Academy of Management Journal* and *Administrative Science Quarterly* (two high-quality, high-impact journals) between January 1976 and September 1996, Locke and Golden-Biddle (1997) concluded that researchers who had successfully demonstrated the novelty or contribution of their research used two rhetorical strategies. First, they 'legitimise' their work by 'constructing intertextual coherence'. They 're-present and organise existing knowledge so as to configure a context for contribution' (p. 1029). Second, they 'subvert' or 'problematise' the existing literature. They do so to show that opportunities exist for contributions to knowledge. One way in which the novelty of a theory can be assessed *ex ante*, therefore, is to evaluate how well its proponents enact Locke and Golden-Biddle's two strategies.

Parsimony

A theory is parsimonious when it achieves good levels of predictive and explanatory power in relation to its focal phenomena using a small number of constructs, associations and boundary conditions. What constitutes a 'small number' is in the eyes of the beholder. Nonetheless, Miller's (1956) classic paper on the 'magical number seven, plus or minus two' suggests some guidelines. Humans appear able to manipulate about seven 'chunks' of information in short-

term memory. In this light, one might predict that researchers would deem a theory to be parsimonious if it has no more than about seven constructs, seven associations and seven boundary conditions (and perhaps the desired number in each case is less than seven).

In building a theory, researchers are often tempted to include *more* constructs, *more* associations and *more* boundary conditions in an attempt to capture the 'richness' of the phenomena that they are seeking to predict or explain (and my experience is that the inclusion of more constructs, associations and boundary conditions is often a frequent request made by the reviewers of journal papers!).

Parsimony dictates, however, that some constructs, associations and boundary conditions must be *omitted* from a theory. In choosing constructs to omit, those whose instances have little variation in their values (states) are likely candidates. In choosing associations to omit, those where few instances of constructs are related are likely candidates. In choosing boundary conditions to omit, those where only a small number of states and events (pertaining to a construct instance) fall outside the boundary condition are likely candidates to omit.

Often, a trade-off must be made between parsimony and a theory's predictive and explanatory power. As the number of constructs, associations and boundary conditions in a theory increases, the theory might be better able to predict and explain the focal phenomena. At some point, however, users of the theory will deem it to be too complex. The goal is to achieve high levels of prediction and explanation with a small number of theoretical components (Occam's razor).

Level

Some theories cover a very narrow, constrained set of phenomena. They are often called 'micro-level' theories. On the one hand, a micro-level theory's constructs, associations and boundaries might be defined precisely. Moreover, its explanatory and predictive powers might be high in relation to the phenomena it covers. Because of the limited range of phenomena it covers, however, it runs the risk that it will be deemed uninteresting.

Some theories cover a broad range of phenomena. They are often called 'macro-level' theories. In some ways, a macro-level theory might be compelling because of the overall insights it provides into many phenomena. Often, however, its constructs and associations are defined imprecisely. Moreover, its explanatory and predictive powers in relation to the more specific phenomena that are a researcher's focus are limited. It runs the risk that it will become discredited because it ends up being a 'theory of everything' in a discipline.

Merton (1957) argues the primary theories used by a discipline ought to be 'middle-range' (or 'meso-level') theories. On the one hand, such theories avoid

'narrow empiricism'. On the other hand, they avoid being so general in their coverage that it is difficult, if not impossible, to test them empirically. Meso-level theories often have value because they link the micro-level world and macro-level world in a discipline.

In spite of the wide acceptance of Merton's idea within many disciplines, the precise meaning of 'middle-range theories' remains problematic (Boudon, 1991). Whether a theory is at an appropriate level is a matter of judgment. Moreover, a level that is too high or too low in one discipline might be an appropriate level in another discipline. Nonetheless, in the context of their discipline, researchers make judgments about whether a theory is at an appropriate level—whether it is too specific or too broad to be interesting and/or useful.

Falsifiability

Most, if not all, theories cannot be *proven* via empirical tests, because it is impossible to test the theory under: a) *all* combinations of values that its *constructs* might assume; and b) *all* combinations of values that its *boundary conditions* might assume. Instead, *support* for a theory grows when its powers of prediction and explanation remain robust across different tests of the theory. If the theory has been articulated clearly, these tests can be designed strategically. They can be used to examine conditions that researchers believe are most likely to lead to the theory being *falsified* (failing the empirical test) rather than supported (Doty and Glick, 1994; Popper, 1990).

To be capable of falsifying a theory, researchers must be capable of generating sufficiently precise predictions about the focal phenomena so they can undertake reasonably exact empirical tests of the theory. If the predictions they are able to generate are so vague that the status of empirical tests they undertake always remains problematic or alternatively the empirical outcomes can always be 'finessed' (explained) using the theory, the value of the theory is undermined.

Using the Evaluation Framework: An example

To show how the evaluation framework I have proposed above can be used, I examine in the subsections below the paper by Griffith et al. (2003), which examines 'the dynamics of knowledge development and transfer in more and less virtual teams' (p. 265). This paper is one of several published in a special issue of the *MIS Quarterly* on the topic of 'Redefining the Organizational Roles of Information Technology in the Information Age'. The stated purpose of the special issue was to 'stimulate *significant and innovative theoretical thought* in response to the dramatic changes that had occurred in the 1990s regarding

information technology and the transformational ways in which information technology was being applied to enable new forms of organisations and markets' (Zmud, 2003, p. 195; emphasis in original).

In spite of the special issue's focus on 'significant and innovative theoretical thought', it is unclear whether Griffith et al. (2003) are seeking to articulate a *theory* or a *model* in their paper (recall, I argue above that not all models are theories and that the latter are models that possess specific attributes). On the one hand, they state their paper 'advances theory' (p. 265). Moreover, they articulate a number of propositions, which suggests that their focus is theory building. On the other hand, they present a 'stylised *model*' (my emphasis) of how 'individual and social knowledge…transfers among individuals…becomes available to the members of the team' (pp. 268–9). Moreover, they use the term 'model' frequently throughout their paper. Nonetheless, the legend they give for their Table 1 of their paper (p. 281) is 'Operationalisation of Constructs to Test the *Theoretical Model*' (my emphasis), which suggests their model is indeed meant to be a theory. In any event, for the sake of illustrating how the evaluation framework might be used, I have assumed that Griffith et al. are seeking to present a *theory* of virtualness and knowledge in teams.

Parts

In this subsection, I evaluate the constructs, associations and boundary that Griffith et al. (2003) employ in their theory. Specifically, I focus on how rigorously each is specified in their paper. In this regard, while the evaluation framework I have proposed above pinpoints those parts of Griffith et al.'s theory that need to be assessed, readers of their theory still need to make judgments about how rigorously each part has been expressed (and accordingly my evaluation below reflects my own judgments). Even where judgments about rigour differ, however, the evaluation framework provides a way for researchers to structure their discourse about the quality of a theory's components.

Constructs

Griffith et al. present their constructs at four places in their paper. First, they are shown in Figure 2 (their 'stylised model') of their paper (p. 269). Second, they can be gleaned from the 19 propositions they state in their paper (pp. 271–8). Third, in Table 1 of their paper (p. 281), they state that they 'catalogue the constructs and assessments necessary to test our model' (p. 280). Fourth, they discuss specific constructs at various places in the text of their paper.

A first problem with Griffith et al.'s articulation of their constructs is that inconsistencies exist among those they show in Figure 2 of their paper, those embedded within their propositions and those they list in Table 1 of their paper.

In this regard, Figure 2 of their paper appears to show 17 constructs they employ in their theory. In my reading of their propositions, however, I can identify 29 different constructs that they employ in their theory (see my Table 1.1 below). Yet Table 1 of their paper shows only 15 constructs that must be subject to 'assessments'.

Table 1.1 Constructs in Griffith et al.'s Theory of Virtualness and Knowledge in Teams

No.	Thing	Attribute
1	Team	Level of virtualness
2	Team	Level of transformation of implicit knowledge into explicit knowledge
3	Team	Level of access to extant explicit knowledge
4	Team	Level of proactive effort made to verbalise rules, terminology and descriptions
5	Team member	Level of tacit knowledge acquired from collocated sources transferred to team
6	Team member	Level of tacit knowledge acquired from team-mates
7	Team	Level of difficulty in forming collective knowledge
8	Team	Level of experienced richness of communication
9	Team	Level of collective knowledge
10	Team	Level of collective knowledge accessible via technological tools
11	Team	Likelihood of enacting an independent approach to tasks
12	Team	Level of shared understanding of tasks
13	Team	Level of access to and appropriation of tools and structures that support highly interdependent work
14	Team	Level of interdependence of work
15	Team	Level of shared knowledge
16	Team member*	Level of transition of potential team knowledge to usable knowledge
17	Team member	Level of absorptive capacity
18	Team member	Level of social interaction limited by virtual work undertaken
19	Team	Level of transition of potential knowledge to usable knowledge
20	Team member	Level of connections to relevant communities of practice
21	Team	Level of access to communities of practice
22	Team	Level of tacit knowledge from members' links to communities of practice disseminated within team
23	Team	Level of transfer of potential team knowledge to usable team knowledge
24	Team	Level of transactional memory
25	Team	Level of transactional memory development
26	Team	Level of virtual work
27	Team	Extent to which technologies of organisational systems are used to support transactional memory development
28	Team	Level of synergy
29	Team	Extent of match between team task and technology use

* indicates it is unclear whether the attribute belongs to 'team' or 'team member'.

Prima facie, some of these inconsistencies appear to represent only *naming* inconsistencies. For instance, Figure 2 of Griffith et al.'s paper shows a construct called 'Individualised Knowledge: Implicit', which is cross-referenced to Proposition 2 (P2) in their paper. Based on the label given to this construct, one might expect it refers to the *level* of implicit knowledge that a member of a virtual team possesses. The focus of Proposition 2, however, is on the *extent to which implicit knowledge can be transferred to explicit knowledge*. These are *not* the same constructs, even though the words 'implicit knowledge' are used in both. Furthermore, the 'assessment' (operationalisation) of the 'Individual Knowledge Types: Implicit' construct in Table 1 of Griffith et al.'s paper does not pertain to the extent to which implicit knowledge can be *transferred* to explicit knowledge (the construct used in P2). Rather, it refers to the 'extent to which individuals rely on…knowledge which could be codified but has been made automatic by practice' (p. 281).

A similar problem exists with other constructs—that is, the meaning that prima facie might be assigned to a construct shown in Figure 2 of Griffith et al.'s paper does not match the construct employed in their propositions. In addition, the construct used in their propositions does not match the construct in Table 1 of their paper.

A second problem with Griffith et al.'s articulation of their constructs (which is to some extent a corollary of the first problem) is that some are defined rigorously (for example, the level of team virtualness and 'individual knowledge types'), but others are not. Moreover, the meaning of some constructs has to be elicited from the text the authors use to articulate and support their propositions. Sometimes the meaning of these constructs is clear; sometimes it is not.

For instance, a construct Griffith et al. use in their Proposition 4b is 'Level of Collective Knowledge Accessible Via Technological Tools'. Earlier in their paper (p. 273), they define collective knowledge reasonably precisely as 'explicit knowledge that has been internalised by the team members'. What they mean by 'technological tools', however, is discussed only somewhat obliquely. Nonetheless, it is not clear which of the following meanings they ascribe to the construct they use in Proposition 4b: a) the *nature* of the collective knowledge formed by more virtual teams means that this knowledge is easier to access via 'technological tools'; or b) more virtual teams have more access to or greater facility with 'technological tools' and thus find it easier to access collective knowledge; or c) both meanings apply to the construct. If the theory is to be rigorously operationalised (tested), the meaning of the construct must be clarified.

One approach that Griffith et al. might have used to clarify the meaning of *all* their constructs is to employ a table similar to Table 1.1 above. In the table, they

could have shown the *things* that underlie each of their constructs (team or team member) and the attributes associated with the things. They also could have provided a rigorous definition of the construct. In the absence of Griffith et al. having defined all their constructs precisely, it is difficult to test their theory empirically. The reason is that valid and reliable measures cannot be devised for constructs that are not defined rigorously. Table 1 of their paper (p. 281) shows a number of constructs for which '[m]easures have to be developed', but valid and reliable measures cannot be developed unless the meaning of each construct is clear.

Associations

Griffith et al. state 19 propositions in their theory. Nine of these propositions manifest a single directional association between two constructs (five positive associations and four negative associations). Two propositions (P5a and P7) manifest two directional associations involving three constructs (one construct is associated with another construct that in turn is associated with another construct). Eight propositions manifest *moderated* associations—in other words, the strength of the directional association between two constructs is moderated by a third construct (an interaction effect is postulated).

Griffith et al.'s use of directional and moderated associations strengthens the potential predictive and explanatory power of their theory. Moreover, while they do not use the terms 'cause' and 'causal' when discussing their propositions, nonetheless causality is implied in the arguments they provide to support many of their propositions. For instance, it seems clear that they believe the existence of virtuality in a team *causes* certain outcomes to occur in relation to how different types of knowledge are transferred among team members. To the extent their propositions imply causality either implicitly or explicitly, the predictive and explanatory power of their theory is enhanced further.

Some arguments provided by Griffith et al. to support some of their associations are rigorous and compelling; however, two factors undermine the rigour of the arguments they use to support other associations. First, as discussed above, some of their constructs are not defined clearly. As a result, the meaning of any associations that employ these constructs will lack clarity. Second, because of the large number of constructs and associations employed in their theory, it is difficult to provide rigorous argumentation in support of all of them. Inevitably, some associations will be better argued than others.

Griffith et al.'s failure to specify all the associations in their theory rigorously undermines researchers' ability to test their theory empirically. In the absence

of each association being articulated rigorously, researchers will lack the understanding they need to be able to evaluate whether the theory's associations hold empirically when they observe the outcome of a test of the theory.

Boundary

Griffith et al. do not use the term 'boundary' within their paper, nor do they have a specific section in their paper that discusses the boundary of the theory they are proposing. Nonetheless, at one point in their paper they indicate that their theory is not applicable to all kinds of virtual teams: 'This model is presented from the perspective of virtual teams where membership is relatively stable, but with members having interaction both within the focal team, as well as with collocated others' (p. 269).

Use of the evaluation framework motivates considerations of whether the theory is constrained in other ways. For instance: does it apply to all kinds of tasks that a virtual team with a relatively stable membership might undertake? Does it apply when the virtual team is made of members having substantial differences in culture? Does it hold throughout all phases of the virtual team's existence? Griffith et al. are silent on such questions. In the absence of their specifying the boundary to their theory clearly, however, researchers might test their theory in an inappropriate context.

Whole

In this subsection, I evaluate Griffith et al.'s (2003) theory as a whole. The evaluation of their theory's emergent attributes is more judgmental than the evaluation of their theory's parts.

Importance

In the introduction to their paper, Griffith et al. provide some clear and compelling reasons why the phenomena they are investigating are important for practice. They point out that the management of teams and knowledge is an important way of creating 'synergies in…resources' and 'increased value' for organisations (p. 266). Moreover, with the emergence and ongoing refinement and development of collaboration technologies and the increasing globalisation of workforces, virtual teams are becoming more prevalent. Thus, the successful operation of virtual teams is now critical to the success of many organisations (for example, Lowry et al., 2010).

From a research perspective, Griffith et al. argue their research potentially provides a foundation for other researchers who wish 'to identify the limiting conditions for effective learning and knowledge transfer across the range of

traditional, hybrid, and virtual teams' (p. 280). It is clear that they have achieved this outcome, because Google Scholar shows their paper has been cited approximately 300 times.

Novelty

Prima facie, it does not appear that Griffith et al.'s paper has been paradigm changing in the sense that it has fundamentally altered the ways in which researchers view phenomena associated with virtual teams and knowledge transfer. In short, they are following a normal-science approach in their research (Kuhn, 1996). Nonetheless, their research can be deemed novel for several other reasons.

First, at the time their paper was published, their theory included a number of constructs that, if not completely new, had received only cursory attention in the extant research literature. For instance, Table 1 of their paper (p. 281) shows several constructs where they note 'measures to be developed'. Table 1.1 above also contains a number of constructs that, to the best of my knowledge, have not been canvassed extensively by researchers (for example, 'level of social interaction limited by virtual work undertaken').

Second, their paper included a number of associations that had received either no or only cursory attention in the research literature that existed at the time they prepared their paper. For instance, based on their analysis of the existing literature, the eight moderated associations they proposed in their theory appear to be new.

Third, the 'package' of constructs and associations included in their theory was novel. While at the time their paper was prepared other researchers might have canvassed *subsets* of the constructs and associations covered by Griffith et al.'s theory, the 'whole' was new. Their theory covered team virtualness and knowledge transfer phenomena in novel and potentially interesting and important ways.

In the context of Locke and Golden-Biddle's (1997) two strategies for demonstrating the contribution to knowledge of a piece of research, Griffith et al. first *construct inter-textual coherence* using Locke and Golden-Biddle's tactic of 'synthesised coherence'—making connections between literatures that historically have been somewhat disjointed (Locke and Golden-Biddle, 1997, pp. 1030–5). They enact Locke and Golden-Biddle's second strategy, *problematising the existing literature*, by using the tactic of 'incompleteness'—that is, showing that the existing literature can be characterised by knowledge gaps or lacunae (Locke and Golden-Biddle, 1997, pp. 1030–5).

Those tactics are manifested in the way they frame the contribution of their paper: 'The model is largely drawn from the extant literature…Our contribution is in combining the results from the prior literature in a way that is amenable to an assessment of the opportunities and challenges presented by considering more and less virtual teams from the perspective of knowledge' (p. 270). In short, Griffith et al. have tacitly followed Locke and Golden-Biddle's recommendations for demonstrating novelty and contribution via the rhetoric used to contextualise a piece of research.

Parsimony

As I indicated above, I believe Griffith et al.'s theory contains

- 29 constructs (rather than 17 constructs, as shown in Figure 2 of their paper, or 15 constructs, as shown in Table 1 of their paper)
- 19 associations (these are shown as propositions in their paper)
- One boundary condition.

Based on a simple count of the number of constructs and associations in Griffith et al.'s theory, it is difficult to conclude their theory is parsimonious. As a result, one might predict that this lack of parsimony would undermine the impact of the theory on other researchers. Interestingly, as indicated above, citation data suggest otherwise. Given the large number of citations of Griffith et al.'s paper, their theory clearly has had an impact on other researchers. Thus, contrary to expectations, lack of parsimony has not undermined the impact of their paper.

Level

In my view, Griffith et al. have articulated a middle-range theory. The range of phenomena that their theory covers is reasonably broad. Thus, they cannot be accused of narrow empiricism. Moreover, while a number of their constructs have yet to be defined rigorously and to be operationalised, it is possible to conceive how ultimately these outcomes might be achieved. In short, their theory is framed at a level that enables it to be employed to generate useful predictions and insights about, and understanding of, their focal phenomena.

Falsifiability

I have argued above that some parts of Griffith et al.'s theory have been articulated clearly and that other parts of the theory have *not* been articulated clearly. Where clarity exists, empirical tests can be undertaken to test the theory. Potentially, the outcomes of these tests will lead researchers to conclude that Griffith et al.'s theory is not supported. That is, the theory can be falsified. For those parts of their theory that are not articulated clearly, however, attempts to falsify the theory are problematic. Empirical tests that produce 'unfavourable' outcomes

might simply mean that researchers have used invalid or unreliable measures for constructs or that they have failed to understand the nature of an association between constructs. They also might have tested the theory in a context that falls outside its boundary.

Conclusions

The framework I have proposed above facilitates an evaluation of the quality of an existing theory. It also informs researchers who are seeking to build a new theory or enhance or adapt an existing theory. As they construct their new or modified theory, they should be mindful of matters they need to address from the perspective of achieving high-quality outcomes in relation to the parts and whole of their theory. In essence, the framework can be used as a set of checkpoints to test the quality of the work they are undertaking.

The framework does not assist, however, in choosing the focal phenomena and the ways these phenomena might be conceived, nor does it assist in choosing a theory's constructs, associations and boundary. To a large extent, these choices remain creative acts that affect, in particular, the quality of the whole—a theory's importance, its novelty, its parsimony, and so on. In the information systems discipline (and in a number of other disciplines), I believe a rich vein of research lies in seeking to better understand the characteristics of those choices that have led to the articulation of high-quality, high-impact theories.

References

Boudon, R. (1991). Review: what middle-range theories are. *Contemporary Sociology, 20*(4), 519–22.

Bunge, M. (1977). *Treatise on Basic Philosophy. Volume 3. Ontology I: The furniture of the world*. Dordrecht, Netherlands: D. Reidel Publishing Company.

Bunge, M. (1979). *Treatise on Basic Philosophy. Volume 4. Ontology II: A world of systems*. Dordrecht, Netherlands: D. Reidel Publishing Company.

Doty, D. H., & Glick, W. H. (1994). Typologies as a unique form of theory building: towards improved understanding and modeling. *Academy of Management Review, 19*(2), 230–51.

Dubin, R. (1978). *Theory Building*. [Rev. edn]. New York, NY: Free Press.

Eisenhardt, K. (1989). Building theories from case study research. *Academy of Management Journal, 14*(4), 532–50.

Gregor, S. (2006). The nature of theory in information systems. *MIS Quarterly, 30*(3), 611–42.

Griffith, T. L., Sawyer, J. E., & Neale, M. A. (2003). Virtualness and knowledge in teams: managing the love triangle of organizations, individuals, and information technology. *MIS Quarterly, 27*(2), 265–87.

Kuhn, T. S. (1996). *The Structure of Scientific Revolutions*. [Third edn]. Chicago, Ill.: University of Chicago Press.

Lewin, K. (1945). The research center for group dynamics at Massachusetts Institute of Technology. *Sociometry, 8*(2), 126–35.

Locke, K., & Golden-Biddle, K. (1997). Constructing opportunities for contribution: structuring intertextual coherence and 'problematizing' in organizational studies. *Academy of Management Journal, 40*(5), 1023–62.

Lowry, P. B., Zhang, D., Zhou, L., & Fu, X. (2010). Effects of culture, social presence, and group composition on trust in technology-supported groups. *Information Systems Journal*, *20*(3), 297–315.

Markus, & Robey, D. (1988). Information technology and organizational change: causal structure in theory and research. *Management Science, 34*(5), 583–98.

Merton, R. K. (1957). *Social Theory and Social Structure*. Glencoe, Ill.: Free Press.

Miller, G. A. (1956). The magical number seven, plus or minus two: some limitations on our capacity for processing information. *Psychological Review, 63*(2), 343–55.

Mone, M. A., & McKinley, W. (1993). The uniqueness value and its consequences for organization studies. *Journal of Management Enquiry, 2*(3), 284–96.

Popper, K. (1990). *The Logic of Scientific Discovery*. London, UK: Unwin Hyman.

Weber, R. (2003). Theoretically speaking. *MIS Quarterly, 27*(3), iii–xi.

Zmud, R. (1998). Editor's comments. *MIS Quarterly, 22*(2), xxix–xxxi.

Zmud, R. (2003). Special issue on redefining the organizational roles of information technology in the information age. *MIS Quarterly, 27*(2), 195.