Learning Architectures in Higher Education

Beyond Communities of Practice

Jonathan Tummons

Learning Architectures
in Higher Education

Also available from Bloomsbury

Learning Communities in Educational Partnerships, Máirín Glenn, Mary Roche, Caitriona McDonagh and Bernie Sullivan
Reflective Teaching in Higher Education, Paul Ashwin

Learning Architectures in Higher Education

Beyond Communities of Practice

Jonathan Tummons

Contents

Acknowledgements

Communities of Practice

Introduction

In this chapter I introduce two linked topics for discussion: communities of practice and learning through legitimate peripheral participation (LPP). I provide a working definition of each of these and briefly outline some of their more problematic aspects: as the book continues, a more sustained exploration will emerge. I also provide a first indication of some of the different theoretical and conceptual approaches that different writers (including myself) have drawn on in order to solve some of the problems that are latent in communities of practice theory. The aim of this first chapter, therefore, is to set the scene for what follows, to provide a quick but nonetheless critically aware overview of our topics, and to begin to open up both communities of practice and LPP to a more detailed investigation.

Introducing communities of practice

I start with a list of stuff that people do. Some of these things are jobs, others are hobbies or pastimes and some are everyday activities that count as work (in the sense that 'work' can be any kind of purposeful and effortful activity) even if they do not attract remuneration. Here is the list: amateur radio operators; athletes; recovering alcoholics; high-school teachers of mathematics; adult learners in a basic skills class; office managers; trainee nurses on placement; butchers; human resources officials; researchers; people who are unemployed; vocational trainers; architects; and higher education (HE) lecturers.

These are some things that people do that have been written about, researched and theorized through reference to *communities of practice*. That is to say, researchers and writers (including me) have written about these (and

many other) more or less discrete or straightforwardly definable groups of people who do particular jobs, have particular hobbies, or otherwise share particular kinds of work or practice in the world, and have *theorized* what these people do, how they talk, what tools or objects they make and use, and, most importantly, how they learn, using theoretical perspectives and concepts derived from a large, and ever-growing, body of work that, for the sake of simplicity, I shall refer to as being based on *communities of practice theory*. So we might find references to a community of practice for 'university-based teacher educators' (Herrington et al., 2008), communities of practice 'in inclusive education' (Mortier et al., 2010) or a community of practice in a 'secondary mathematics classroom' (Goos, 2004).

It is a highly influential theory (if 'theory' is what it is – it might be a 'concept' or a 'framework', or something else entirely) and a popular one as well. It is also a troublesome theory, a theory that is used badly as often as it is used well. But I go on to argue that this is also the case with the use of several other theories in HE research and practice. And it is important to reflect on the fact that it is an attractive theory, which may well explain elements of its popularity. The term 'community of practice' seems to be quite straightforward for people to pick up and use. I have lost count of the number of conference presentations that I have listened to or journal articles that I have read that set out to explore a community of practice, to investigate how one might be established, or to provide an account for how being in one might contribute to education, learning or the transfer of knowledge.

It seems right to say that if a proposed theoretical framework that is based on empirical study can accommodate practices as diverse as some of those that I have listed above, then to go about dismantling that same theoretical framework must be seen as being of questionable value, not to say wisdom: a framework that can be employed so straightforwardly across so many different contexts must be doing something right. But at the same time it seems right to pause for a moment and consider that if we have a theory that can be used to make sense of practices as diverse as those listed above, then does that mean that the theory is being used indiscriminately, or that it is too big or too diffuse? Does it make sense to use the same theoretical framework to explore how apprentice tailors learn their craft, how university lecturers learn to grade essays or how office managers leverage innovation within their organizations? These all seem to me to be rather different from each other. And even if we focus specifically on HE, the range of contexts that has been explored through communities

of practice theory seems remarkable. Recent research has focused on academic disciplines (architecture, education, linguistics, psychology) as well as academic work (assessment moderation, curriculum reform, online learning, staff development). Even just a cursory glance through the literature indicates a considerable variety in the kinds of practices and communities that have been researched and written about.

This variety extends to the meaning that is attached to the term 'community of practice', however, and this is the first complication that we need to address. This complication has been created and then exacerbated by those researchers and writers who have used the term 'community of practice' in more or less robust ways, for example, in describing a group of people as a community of practice but without providing an account as to *why* they are a community of practice, or in focusing so much on the community that the learning that is happening – and it is a theory of learning – gets left to one side. Perhaps a more difficult problem lies in the fact that one of the begetters of the theory has fundamentally changed his definition and understanding of what a 'community of practice' actually is (Farnsworth et al., 2016; Lave and Wenger, 1991; Wenger, 1998, 2000; Wenger et al., 2002; Wenger-Trayner et al., 2014). In order to make sense of these differing and sometimes conflicting ideas, it seems sensible to spend some time thinking about what a community of practice actually is, and how thinking about them can help us to understand learning as a social practice. As well as helping us to focus once again on learning (which is what the theory is all about), such an account can also provide the context for the empirical contribution to the conversation about communities of practice that I wish to make in this book, namely, that through a focus on Etienne Wenger's concept of *learning architecture* we can find creative as well as theoretically robust ways to think about learning and teaching in universities. Learning architecture is a component of Wenger's original (1998) theorization, and it is this iteration of communities of practice theory that this book rests on more broadly (although not without some criticisms and reservations). In order to fully unpack the problems of defining our key terms, it is necessary at different points to consider how the later iterations of communities of practice differ from the original, and place their emphases on different aspects of learning within quite different social and organizational contexts. But I start with a brief account of what I think communities of practice are and how learning as a social practice within communities of practice can be understood, before moving on to consider some of the problems that these ideas bring with them (Tummons, 2012, 2014a).

What are communities of practice?

When introduced by Jean Lave and Wenger in their book *Situated Learning: Legitimate Peripheral Participation*, the term 'community of practice' was left relatively unexplored, only loosely defined as a 'largely intuitive notion' that required further investigation (1991: 42). Arguably, the term was introduced as a by-product of their more sustained analysis of learning as *legitimate peripheral participation* (which was in fact the focus of the book, and to which I shortly return), in order to create some sense of the kinds of cultural and social places where learning might happen. How a community of practice might be identified, described or defined, or questions relating to what the constituent components or characteristics of such communities might be, were only later explored in depth by Wenger (1998). So, what are communities of practice and where might we find them?

Communities of practice are everywhere. We are all members of multiple communities of practice. Some of these communities work in relative isolation, while some overlap with others. Sometimes we actively seek out membership, but at other times we are not even aware that we are members of a particular community. As people in a social world, we engage in all kinds of activities – *practices* – as part of our everyday lives, interacting with other people, sometimes in close proximity and sometimes at a distance or by proxy: at work, at play, with families or with friends. In order to take part in these various practices people come together in *communities* so that they can talk about, share and learn more about them. These *communities of practice* can be found in formal, institutionalized settings and in informal, vernacular ones. Lave and Wenger's examples include tailors, midwives and butchers (1991). Wenger's examples include amateur radio operators, recovering alcoholics and office-based computer users (1998). In some communities, members meet and talk on a regular basis; in others, they meet only infrequently. Some communities have existed for a long time; others are relatively new. Some communities establish and sustain close relations with others, sharing aspects of their practice, while others are relatively self-sufficient. All communities of practice, however, share one characteristic: they are all socially configured spaces that necessarily involve learning as an aspect of membership. That is to say, for any member of any community of practice, learning is an ever-present aspect of participation – even if the learning is unintentional, or the member in question does not realize that they are learning.

In his later book, *Communities of Practice: Learning, Meaning and Identity*, Wenger takes the time to point out that not all communities of people are

'communities of practice'. It is perfectly possible for groups of people to come together in all kinds of ways in the social world without being part of, or forming, a community of practice. Therefore, in order to identify whether a community is in fact a community of practice, Wenger outlined a number of specific structural qualities that form around any such community. Three such qualities, or attributes, need to be looked for and, I would suggest, richly described: *mutual engagement, joint enterprise* and *shared repertoire* (Wenger, 1998: 73–85).

Mutual engagement is the term used by Wenger to refer to all of those ways through which the members of a community of practice interact with each other and do whatever they do. Members of a community might engage with others in a complementary manner or an overlapping manner, depending on the relative competence and positions that they occupy: some people are newbies, and others are more expert; some people pick up and learn new things quite quickly, while others take longer. Mutual engagement can take different forms: in some communities, engagement might always be done on a face-to-face basis; in others, it might involve talking on the telephone as well as face to face, or chatting over Skype, or posting messages on Facebook or Twitter. Mutual engagement might require talking, writing and reading (on screen, on paper), listening, moving or making. Because working together creates differences as well as similarities, mutual engagement is never homogeneous. Members do not have to agree with each other all the time: in fact, change is a constant element of many, if not all, communities of practice. Things can be done, adjusted, argued over, tried differently or spoken about in various ways so long as these are, in the end, reconcilable to the *joint enterprise* of the community of practice.

Joint enterprise refers to the shared work or endeavour of the community of practice – the thing or stuff that the community is about. So long as all the members keep that work in mind, community cohesion can be established and maintained. The joint enterprise of a community might involve a specific activity that entails physical effort, mental effort emotional effort, or – probably more likely – any combination of these. It might involve a small number of people who meet up in the same physical location on a monthly basis, or a larger number of people who communicate through a variety of means. In a way, it does not matter exactly how many people might be involved or how often they meet up or speak with each other: what is important is that they all stick to the job at hand, that they all continue to engage in practice in order to sustain, change or otherwise maintain the joint enterprise of the community.

Doing the work of the community might require people to write things down, to take part in shared and routinized activities, to watch and comment on tasks,

or to make or break different kinds of items, artefacts or bodies of materials. In order to engage in practice, therefore, members draw on the habits, discourses, routines, ways of talking, tools, structures and other artefacts that over time have been created or adopted by a community of practice: the *shared repertoire* of the community. Such structures and artefacts serve a number of functions. They allow the members of a community to make statements about their practice, to express their identities within the community, and they represent the history of mutual engagement within the community. The repertoire can be seen as taking aspects of the practices of a community and turning them into solid forms – a process that is called *reification* (to reify something means to turn a concept or mental construct into a physical thing – for example, abstract notions of justice can be reified into statutes, which can then be written down and passed around). The repertoire can also be seen as reflecting the different ways in which members engage in practice, so members draw on the repertoire of the community in differential ways as they learn. Newcomers probably will not know how all the tools or routines work at first, and so part of their learning within the community will include coming to know how the artefacts of the community work and can be used.

This three-part model based on *mutual engagement, joint enterprise* and *shared repertoire* seems, at first look, to provide us with a helpful – and straightforward-to-operationalize – framework to ascertain the extent to which a collection of people, all engaged in some kinds of action, constitutes a community of practice. Indeed, it seems to be sufficiently straightforward to beg the question as to why it is not used more often in academic literature when people describe something as a community of practice (and then go on to drop in a reference to Wenger (1998)). I earlier suggested that the use of this three-part model would simply require the researcher to look for and richly describe each element: this is merely a reflection of the anthropological and ethnographic roots of the landmark books written by Lave and Wenger, and Wenger. But all too often we find collections of people or groups of practices described as being a 'community of practice', without any serious attempt to establish why they are so. In circumstances such as these, we find the phenomenon of poor or uncritical use of theory (Thomas, 2007). At the same time, we find a number of theoretical problems and questions that Lave and Wenger, and Wenger have opened up but failed to provide full answers to, and which other people have followed up instead. What about the power relations between members as they engage in practice? How diffuse can joint enterprise be while still being coherent enough to sustain a community? What can or cannot be done with the tools, artefacts, rules and routines

of a community? What do we mean when we talk about 'community'? What is 'practice'? These are indeed difficult problems to consider: on the one hand, we have a theory being poorly considered, badly used and misunderstood; and on the other, a theory that does not necessarily stand up to rigorous scrutiny and that raises questions as to its coherence and ability to address the problems that it raises. And in both cases, it is something to do with the use of theory that needs to be addressed.

Interlude: Theory use in higher education research

Issues surrounding the variable use of theory in HE research are well established in the literature (Tight, 2004), and I briefly touch on two key themes here (but I return to these from time to time in later chapters). The first problem to consider is the definition of 'theory', a word that is used in different ways. For James Paul Gee, a theory is a body of generalizations, which can be drawn together to offer explanations and descriptions of the phenomena being researched, which in turn inform peoples' beliefs about things (1996: 16). Malcolm Tight has argued that theories are suppositions that explain something, or seek to explain it, and posits theory as the ability to explain or understand the findings of research within a conceptual framework (2004: 399). Paul Ashwin, likewise, has positioned theory as informing the conceptualization of research: the framing of research questions, the analysis of the data that is created through the research process and understanding the significance of the findings that are drawn. At the same time he warns against using theory to structure research in such a way that the research simply consists of a tautological restatement of the theory in question (2009: 133). Our problem, simply put, is that 'theory' is, all too often, poorly defined and poorly operationalized (Thomas, 2007).

If the first problem is defining theory, then the second problem is how it is used and written about. Here I consider the second problem in two ways. The first rests in the relationship between theory and research. Martyn Hammersley (2008) has argued that one of the failings of qualitative research during the last five decades and more (my own research is undoubtedly 'qualitative', but I reject the spurious divide between qualitative and quantitative research) is the failure of qualitative researchers to develop and then test theory in a systematic manner, a failure that he links to the broader issue of generalizability (and, specifically, the lack thereof) in qualitative research. The second way in which I want to

consider this problem is through thinking about the ways in which some people write about and cite from or refer to theory that sees theory treated as a veneer, a layer of sometimes needlessly complex language, sometimes dropped into an empirical study with relatively little thought as to its applicability or relevance. This has resulted in what Gary Thomas has described as a use of theory now superseded by an excess of *theory talk*, erroneously used to claim 'epistemological legitimacy and explanatory commentary' (2007: 85).

Back to communities of practice

There is not much that can be done with those studies that simply throw a 'community of practice' label onto some empirical work and hope for the best, other than to highlight these as examples of theory talk, rather than research that rests on a robust and critical reading of communities of practice. Fortunately, our other problem can be addressed in a more comprehensive manner: many excellent studies have sought to explore what might be termed the 'problematic elements' of communities of practice theory and these draw on other theoretical perspectives seen as being compatible with the work done by Lave and Wenger, and Wenger. The analysis of power relations within communities of practice has been explored from a Foucauldian perspective, and from the perspective of actor-network theory (Fox, 2000, 2005). The strength – or otherwise – of a community has been explored through network analysis (Jewson, 2007). The creation and use of artefacts, specifically text-based artefacts, have been expanded through the theoretical insights of literacy studies (Barton and Hamilton, 2005). Later, I discuss these and other ways of moving within and beyond communities of practice in more detail. For now, however, I want to take a step back from communities of practice and focus on learning.

Learning: Legitimate peripheral participation in communities of practice

Lave and Wenger's book, published in 1991, in which the 'community of practice' is first, and somewhat loosely, defined, is not about communities of practice at all; rather, it is about learning. Specifically, it is about a particular model of learning that they called *legitimate peripheral participation*, and that they derived from ethnographic research across a range of contexts (some of which was only

published much later (Lave, 2011)). Lave and Wenger were interested in generating a descriptor of learning that foregrounded engagement in social practice. Learning, they argued, is an integral element of social practice, a practice that involves the whole person, rather than being 'simply' a cognitive function. They described learning as an improvised practice, not relying on any formal intention to teach or otherwise construct a curriculum. Instead, drawing on a broader social ontology and epistemology as well as more specifically on apprenticeship models of learning within a Vygotskian tradition, they suggested that learning happens as a consequence of participation in any social practice, but with two necessary preconditions: first, that the participation had to be legitimate – that is, authentic and meaningful and not a simulation or simulacrum of the things to be learned; and second, that the participation had to be peripheral – that is, at a level appropriate to the newcomer, or apprentice, at the start of their practice. If some of these ways of thinking about learning seem perhaps to be lacking in specificity or in authoritative detail, then that is because they are (see also Hughes et al., 2007: 4), and both the central concept of learning through LPP and the associated concept of the community of practice as the social place where learning happens and is afforded to people (only briefly touched on by Lave and Wenger) remain in need of more thorough explication.

The vocabulary of newcomers and apprentices, and of old-timers and masters, highlights quite straightforwardly the ancestral traces of different models of apprenticeship learning that underpin Lave and Wenger's work. At the same time, the social ontology and epistemology that their work rests on can be traced to a body of wider research work that has sought more broadly to shift theories of learning – in relation to children as well as adults, and formal as well as informal educational structures or contexts – from a psychological to a sociocultural perspective (Brown et al., 1989; Chaiklin and Lave, 1996; Cobb, 1994; Lave, 1988; Rogoff, 1990). For Lave and Wenger, learning happens when people participate in practice: indeed, if people are participating in practice, then learning cannot *not* happen. As a sociocultural phenomenon, learning through LPP involves and entails changes to the whole person, and how she or he acts and moves within the social world. Learning changes how people think, act and speak: it changes people's *identities* within their community. Consequently, as members become more expert in the practice of the community, they draw on, employ and even enhance the repertoire, tools and artefacts of the community in an increasingly fluent and expert manner. Their participation, within the community, becomes more full.

Alongside an extensive literature that seeks to explore or unpack the concept of the community of practice, there is a somewhat smaller body of literature that

seeks to critique Lave and Wenger's conceptualization of learning as LPP. In a way this is not surprising: books and articles about different theories or frameworks of learning are plentiful, and if we accept that communities of practice rest on a sociocultural model of learning, then it is a straightforward task to conflate an exploration of communities of practice with an exploration of social practice accounts of learning more widely (Illeris, 2007). That is to say, we can look beyond communities of practice literature for insight into the processes of learning that take place within a community of practice, including those aspects of learning that Lave and Wenger, and Wenger, do not explicate as fully as we might like or require.

I return to these themes in more detail in later chapters, but it is worth mentioning a few key questions at this time (although these have to wait before we can answer them more fully). Lave and Wenger (1991) assume that peripheral participation always leads to full participation: but what if full participation is not the intended goal of the participant, or is kept – perhaps deliberately, perhaps accidentally – out of the participant's reach (Lemke, 1997)? Is the learning of full members the 'same' as for new members of a community (Fuller et al., 2005)? Lave and Wenger (1991) refute any sense that learning can be easily divided up into 'formal' and 'informal' practices, but might it be the case that the context in which the learning is happening (and context is central to a social theory of learning) makes a difference (Boud and Middleton, 2003)? Or is there a more fundamental problem with the theory of LPP, namely, that as a theory it explains what is done and not what is learned (Edwards, 2005)?

However, perhaps the most worrisome aspect of their theory – at least, as far as this book is concerned, as well as many other books and articles – is their stark refutation of formal pedagogy, in terms of curriculum, of professors, of a language of instruction, even of institutions. For Lave and Wenger, the organization of any formal educational institution (they refer to 'schooling', but the point is straightforward to generalize (Lave and Wenger, 1991: 40)) rests on the very kind of individual, cognitive, psychological model of learning that they, in their rich and ethnographic study of different forms of apprenticeship, are seeking to move away from. LPP is not a classroom strategy. If there is to be a curriculum, then it should be a *learning curriculum* consisting of authentic resources viewed from the learners' perspectives, and not a *teaching curriculum* designed by others, which would serve only to limit rather than to expand opportunities for authentic participation (ibid., 93–98). And there is to be no special discourse 'corresponding to the . . . lecturing of college professors' (ibid., 108), designed to be a vehicle for instruction.

Elsewhere, I have referred to this as the *pedagogy problem* (Tummons, 2014a), and it is a problem that needs to be addressed – arguably, more extensively than is currently the case. Those articles and books that focus on what we can conveniently term 'informal contexts for learning' need not worry. But, surely, for all those studies that use Lave and Wenger, and Wenger to underpin studies of formal educational structures – whether these are adult basic skills sessions, elementary mathematics classes or university-level business studies programmes – the pedagogy problem needs to be considered? Does the lack of such a consideration speak to the broader problem of theory use that I have already referred to? In his later book, Wenger goes on to provide a concept that can allow us to address this problem: the concept of *learning architecture*, a concept that has been used only infrequently by other writers. I shall, of course, be returning to learning architecture later in this book. For now, it is sufficient to recognize that within Wenger's work there are ways of thinking about formal as well as informal education and training: they simply need – and deserve – to be read, critiqued and applied more frequently than is currently the case.

One community or many?

One of the key tenets of many theories of learning is *transfer*: that is, if we learn something in one place or about one particular thing, then we can transfer that learning to a different, perhaps new, setting, context or application. If we subscribe to LPP specifically, or to a broader sociocultural turn in theorizing learning more generally, then we can make sense of this problem in terms of the whole person. That is, if we understand learning as concerning the whole person as a social actor in the world and not simply as a process of individualized internalization of bodies or schema of knowledge, then we need to think about transfer in terms of people moving within and across different social spaces, rather than thinking about how an individual person might restructure decontextualized knowledge as a process of individual cognition. Fortunately, communities of practice do not exist in isolation. Although communities can be relatively self-sufficient, some of them establish and sustain close relations with others, and might even share aspects of their practice. Indeed, the practice of one community may be influenced by the practice of another. In order to explain how practices, artefacts or even people from one community might be able to move up, down or across into other communities, carrying meanings