

# Social dynamics in user-generated content platforms

Manuel Arriaga

# Index

1. Introduction .....	4
1.1 Overview.....	10
1.2 Brief notes on user-generated content .....	12
1.3 Brief note on the use of “an individual's field” .....	14
2. Literature review .....	14
2.1 Literature on social tagging .....	15
2.1.1 Social bookmarking as a resource discovery tool .....	15
2.1.2 Taxonomies of social tagging platforms.....	18
2.1.3 Motivation and intended goals of users.....	19
2.1.4 Individual-level vocabulary and its development.....	22
2.1.5 Larger-scale/emergent patterns in social tagging platforms .....	24
2.2 Literature on social networking websites.....	28
2.2.1 Impression management and identity.....	28
2.2.2 Network aspects.....	29
2.3 Literature on online reviews.....	31
2.4 How the present work relates to previous research.....	32
2.5 Why Bourdieu's concept of field.....	33
3. Theory.....	35
3.1 Bourdieu's concept of field.....	37
3.2 The concept of online cultural field.....	42
3.3 Active characteristics in an OCF.....	44
3.3.1 Criteria for the selection of active characteristics.....	44

3.3.2 Definition of active characteristics.....	46
4. Proposed model.....	53
5. Empirical work to be done and methods.....	56
5.1 Choice of sites and data collection.....	56
5.2 Identification of a field, its participants and its contributions.....	58
5.3 Operationalization of constructs.....	59
5.4 Testing relations between constructs.....	61
6. Expected contributions.....	61

# 1. Introduction

User-generated content has been the focus of considerable attention on the part of the media<sup>1</sup> and commentators<sup>2</sup> over the recent years. Since 2004, a series of Internet platforms has emerged which allow users to easily upload, organize, share and access digital content. Later considered paradigmatic cases of the so-called “web 2.0”, websites such as Flickr, Del.icio.us, Youtube, Blogger and MySpace thus became vast repositories of photographs, bookmarks, videos, blog entries and personal profiles contributed by their users. The expression “user-generated content” reflects the most salient feature of these websites: that *none* of their content was being produced by paid professionals hired for that purpose (as had traditionally been the case both on- as well as offline), but that instead their users were simultaneously playing the roles of producers *as well as* consumers of the content on the site.

The enormous amounts of resources accumulated on these sites, combined with the fact that they typically allowed users to upload material in a direct, totally unmediated way<sup>3</sup>, introduced two important difficulties. First, it was not clear how to make the collection of resources *explorable* by the users. Search was not a solution for two reasons: (i) several of these sites dealt with non-text resources (photos, video, URLs) for which authoritative metadata was unavailable; and (ii) unlike browsing, search does not easily lend itself to recreational, not-so-targeted exploration (an activity that one could reasonably expect to be important on these platforms). Second, the sites also did not have a way to rank

---

1 Consider, for illustration purposes, *Time* magazine's nomination of “you” – the three letters superimposed on the screen of an Apple iMac – as its “person of the year” for 2007.

2 Clay Shirky, David Weinberger and Andrew Keen are just three among many other names who gained significant (or at least *greater*) prominence through their engagement in discussions related to user-generated content.

3 Unlike, e.g., the traditional mechanism for incorporating user-generated content in more traditional media: the “submission to the editor”.

or filter resources according to their interestingness. The problem, again, was that of lack of metadata.

Two hallmark features of web 2.0 platforms – namely, inter-user links and collaborative tagging – are best understood as ways to make these collections of user-generated content browsable: both provided platform owners with nearly zero-cost means to impose a structure (or, more correctly, two structures: social and topical) on their repositories and thus make them navigable. Additionally, they created valuable metadata for use in filtering mechanisms: e.g., a resource “tagged by 23,894 users” is more likely to be of general appeal than one which has been tagged by a single individual.

The significance of inter-user links and collaborative tagging, however, is not limited to their contributions to the navigability of online content. More fundamentally, they matter because they allow behavior on user-generated content platforms to go *beyond the individual* production and uploading of resources, where each individual uses the platform as little more than a (publically available) extension of her computer's hard-drive. It is due to these two features that users are not merely “generating content”: they are also actively engaged in characterizing, organizing and valuating the contributions *made by others*.

Yet, even when made aware of the fundamentally social nature of the processes taking place on these platforms, many remain unconvinced. By comparison to the “chatrooms” and discussion fora which had previously dominated both popular as well as academic perceptions of online sociability, on these new websites there is little or no interaction<sup>4</sup>. As a result, individuals come across as mere curators of

---

4 The term interaction is used here to signify “[a] process that involves **reciprocal** stimulation and response between two or more individuals” (APA Dictionary of Psychology 2007; emphasis added). The Penguin Dictionary of Psychology (2001) concurs on reciprocity being a necessary characteristic of interaction: “in social interaction the behavior of one acts as a stimulus for the behavior of another *and vice-versa*” (emphasis added). On most user-generated content platforms, users act upon resources contributed by other individuals (e.g., by ranking or annotating a picture uploaded by another user), yet evidently there is no reciprocal link between them.

personal online spaces whose content is, rather unexpectedly, made accessible to online strangers.

To view user-generated content platforms as an agglomeration of personal online spaces is to misunderstand them. Three different factors play a role in hindering our vision of what happens on such platforms. First, we are blinded by our traditional ideas of sociability and what constitutes interaction<sup>5</sup>. Second, we are only able to see *things* and read/hear *utterances*: visiting (and later thinking about) these websites, our mind is captured by the shininess of the millions of resources they hold and what was said/written by individual participants (e.g., in a comment attached to a picture on Flickr). The *relations* between resources, individuals and the factors which make resources and individuals relate to each other (e.g., an interest in a particular topic or activity) are invisible; we are thus misled into believing they are also irrelevant. Third, we (still) pay too much attention to issues of motivation. This insistence on asking “why does someone publish this?” when confronted, e.g., with a carefully curated collection of photographs on Flickr or a revealing MySpace profile has its roots in two distinct mental habits. First, there is the incongruence between the economicist view of rational, self-interested agency and the unrewarded exertion of effort to produce, organize and share content. Second, we have not yet overcome the “private vs. public” dichotomy and get stuck on the fact that millions of individuals are revealing abundant information about their lives to the world at large in a way that invites us to take up the role of *voyeurs* (from a comfortable, safe distance). Both these factors conspire to focus our attention on the individual when thinking about what happens on user-generated content platforms<sup>6</sup>.

In my dissertation I propose to gain a better understanding of the social dynamics on user-generated content platforms by looking *away* from the individual and instead focusing on the relations between

---

5 See previous footnote.

6 Arguably, this is why *Time* chose “you”, rather than “we” or some other plural form, as person of the year.

people, interests and resources.

More specifically, I will study how large online groups engage in the production, dissemination and valuation of user-generated content. I will draw on sociological theory (namely, Bourdieu's notion of field) and test the inducted model on two large datasets obtained from user-generated content platforms.

The central concept in my work is that of an "online cultural field" (OCF). An online cultural field is a group of individuals, and the IT-supported relations among them, who share a cultural affinity and who, by engaging in production, dissemination and valuation of user-generated content, have an influence on the trajectory of other individuals in that field. OCFs occur, for example, on social bookmarking platforms, among bloggers and on social networking websites<sup>7</sup>. One can get an intuitive grasp of the range spanned by these social forms by thinking of them as "more than an Amazon.com co-purchase network, less than an online community" (Provost, personal communication)<sup>8</sup>.

Central to understanding both what OCFs are as well as the motivation for their study is to contrast them with the most often studied online social phenomenon, the online community. An online community is "a large collectivity of voluntary members whose primary goal is member and collective welfare, whose members share a common interest, experience or conviction and positive regard for other members, and who interact with one another and contribute to the collectivity primarily over the Net" (Sproull and Arriaga 2007). Online cultural fields differ from online communities in three

---

7 The reader should keep in mind that a user-generated content platform (such as, e.g., Del.icio.us or Youtube) is **not** an online cultural field. There are multiple online cultural fields on each of those platforms. E.g., on Del.icio.us there are fields centered on programming, design, photography, etc. It is interesting to note that on some platforms (e.g., Del.icio.us and Slashdot) online cultural fields are merely **implicit** and can only be observed (indirectly) by means of identifying commonalities in user behavior, while other platforms support **explicit** affiliation of a user with the online cultural fields she wishes to participate in (e.g., by means of "groups" on Flickr ).

8 It is important to clarify that typically *several* online cultural fields will coexist on an online platform. E.g., the social bookmarking website Del.icio.us is not itself an online cultural field; instead, it is an online social platform which supports a variety of distinct fields.

significant ways: (i) they are not dependent on a notion of group welfare; (ii) their members do not exhibit positive regard for other members; and (iii) fields are not based on interaction among participants<sup>9,10</sup>. By comparison, OCFs constitute a broader (their definition is less stringent), lower-cost (participation entails neither the establishment and maintenance of social ties nor an active, deliberate decision to actually participate) and more strongly technology-mediated (the role of technology is not confined to conveying human-redacted messages) social form.

Although not always in a direct manner, the motivation from my work follows from this comparison. I will first address why I think this work is relevant, and then move on to why I believe it to be interesting.

First, there exist a number of online social phenomena which have been largely ignored by researchers. We will not achieve an adequate understanding of online groups while we continue neglecting these and continue focusing our attention on the more intensely "social" and interactive online forms. This lack of scholarly attention can be understood in part as due to the novelty of such groups themselves (social bookmarking, blogging and social networking platforms did not exist as large-scale phenomena four years ago) and in part as a consequence of the history of research on online behavior (with its traditional comparisons employing face-to-face interaction as the baseline, some degree of similarity – or at least comparability – to face-to-face phenomena was implicitly defined as a criterion of interestingness for online phenomena).

Second, there are reasons to believe that OCFs have the potential to become a widespread

---

9 One might add that “the defining feature of a community is a 'self-consciousness' on the part of each member that the group is a social unit and that he or she shares group identification with the others” (Penguin dictionary of psychology, 2001; emphasis added). By this account, too, online cultural fields are not (online) communities.

10 Another, earlier definition of online community produced by a multidisciplinary group of researchers leads to the same conclusion. Whittaker et al (1997: 137) stress the importance of frequent “intense interactions, strong emotional ties and shared activities occurring between participants” for an online collectivity to be considered an online community.

phenomenon. First (and as mentioned above), OCFs demand low effort from their participants, thus presenting few barriers to participation. Second, and because the definition of an OCF is not modeled on – and thus tightly bound to – a preexisting offline phenomenon, the "space of potential OCFs" is very large: technological innovators have many possibilities to explore in bringing about multiple, diverse platforms supporting OCFs. (By comparison, technological advances can only change so much of online communities if they are to remain online communities.) Third, the current wave of investment into OCF-supporting platforms (e.g., Yahoo has bought Del.icio.us, Murdoch acquired MySpace and Google owns both Blogger as well as YouTube) indicates that their use might soon become even more disseminated.

Third, and building on this last point, there is growing interest in the question of how to derive business value from the existence of these looser, more fluid online social forms. Related work on how to "build" and maintain online communities has been done, but we know nothing about its applicability to OCFs. A better understanding of this phenomenon, one hopes, will help practitioners succeed in building platforms and running businesses based on them.

As a researcher, OCFs strike me as interesting for two additional, related motives (which might or might not add to its general relevance).

First, and as alluded to earlier in this introduction, OCFs constitute an online social phenomenon which lacks an easy, direct counterpart in the offline world. Studying them will, I believe, prove to be a refreshing and important exercise in our research into online social phenomena, which so often focuses on how they relate to some corresponding offline reality.

Second, I find the study of OCFs particularly exciting because it will offer us an opportunity to conceptualize the online social use of technology in a way other than as a vehicle for the exchange of

messages and/or the dissemination of documents. Here we have technology (i) engaged in a much more fundamental way with the establishment of relations among participants and (ii) blurring the border between private and public action<sup>11</sup>. I believe interesting theoretical insights might be derived along both these lines.

## **1.1 Overview**

My work has two central parts. First, I develop a theoretical framework which will allow us to do richer research on these new phenomena by providing a dynamic, (more) comprehensive model of social dynamics on user-generated content platforms. Currently, the majority of academic work on Web 2.0 phenomena either focuses on isolated socio-psychological aspects which were already present (and intensely studied) in the earlier literature on online communities – e.g., motivation for participation, presentation of self, communication behavior and attitudes towards other participants – or consists of atheoretical statistical and graph-analytical explorations of specific online populations. Second, I will then use this framework to study the groups on two large online social tagging and bookmarking platforms, thus illustrating how this analytical lens can bring new insights even to the study of phenomena which have previously received significant attention from the research community.

The structure of this dissertation proposal is as follows. First, the next section presents a review of the relevant literature. By covering earlier work on collaborative tagging and social networking websites, I show how fragmented our understanding of user-generated content is and that we lack a unified, integrative framework for thinking about the dynamics of its production and valuation. I will also very briefly highlight major topics in the study of online communities. Second, in the “theory” section I then introduce Bourdieu's work and proceed to develop, using that as a basis, the concept of an online field

---

<sup>11</sup> This concept, which I call "technological amplification", will hopefully be the subject of future work.

together with a model of its internal workings. Third, in the “empirical work to be done” section I describe the datasets which will be used to test that model and how I will operationalize the different constructs in the two datasets. Finally, this work concludes with a section on expected contributions.

The remaining paragraphs in this overview will provide a sketch of the “theory” and “empirical work to be done” sections in order to make it easier for the reader to place this work in the context of previous studies of user-generated content (while reading the literature review which immediately follows).

Bourdieu's theory of fields is especially useful as the theoretical background for studies of the social phenomena surrounding user-generated content. In his work, a field is a space in which individual agents are positioned according to their possession of different types of "capital" and who are, without necessarily interacting among themselves, permanently influencing the trajectory of other individuals in that space. However, some key aspects of Bourdieu's framework – in particular, its emphasis on social struggle between agents, its three types of capital (economic, social and cultural) and the role of distinction in explaining social dynamics – might make its direct application to the online realm problematic.

A fundamental part in the development of the concept of online cultural field will be the identification and operationalization of "capital-equivalent" constructs. In particular, and by analogy to Bourdieu's focus on forms of capital, I am interested in developing a set of "active characteristics" which will simultaneously condense information about an individual's past behavior as well as give indication of which role(s) that user can perform in the field. The hope is that such constructs will be useful for future discussions of online phenomena, in addition to the constructs that our discipline inherited from psychology (e.g., motivation, communication behavior, etc) and the study of social networks (e.g., network centrality, in- and out-degree, etc).

I then organize these “active characteristics” into a dynamic model of the production and valuation of user-generated content within online cultural fields. This model captures the fundamental ideas of accumulation, preservation and conversion between “forms of capital” which one finds in Bourdieu's theory of fields.

This analytical framework will later be used for my empirical study of two large collaborative tagging platforms. With the cross-sectional data I have collected, I am restricted to reporting correlations among these constructs without being able to effectively test the relations advanced in the model. Access to other data sources, or more skillful inferences from the existing data, will hopefully enable me, at a later point, to conduct more interesting statistical analyses which will probe deeper into the nature of the links between those constructs.

## ***1.2 Brief notes on user-generated content***

Throughout this work I will use the term “user-generated content” (UGC) to refer to the online resources contributed by the users of a web platform which does not discriminate<sup>12</sup> between producers and consumers of the said content.

An important, though only rarely addressed<sup>13</sup>, property of user-generated content (as it is commonly understood) is that users engage with it in a *non-professional* manner. I.e., the vast majority of users come onto these platforms and engage in contributing, consuming and valuating content as activities they choose to conduct outside the scope of their professional lives. (One would not typically think of

---

12 On the aggregate level. Nearly all platforms retain a link between a resource and the user who contributed it; what is meant is that users are not segmented into “contributor” and “reader”/“consumer” accounts with different privileges.

13 See Andrew Keen's book-slash-manifesto “The Cult of the Amateur: How the Internet is Killing Our Culture” for an exception.

scientific papers, e.g., which match the letter of the definition in the previous paragraph, as an instance of user-generated content.)

Though the current work is an attempt at a more general perspective, one should be aware that not all user-generated content is born equal. Although we tend to speak of user-generated content as if it were an homogeneous entity, a moment's thought is enough to readily identify the multiple and markedly distinct forms it can take. Amazon.com product lists, blog posts, Wikipedia articles, photos on Flickr, videos on Youtube and collections of bookmarks on Del.icio.us are all instances of user-generated content. The plan of work proposed in this document consciously incurs the risks inherent in “abstracting away” from possibly fundamental differences in an attempt to unearth some common structure underneath such apparent diversity. (What we are missing by bundling them all together under the rubric of "user-generated content" is an issue I intend to address in my future research.)

For my purposes in this proposal it suffices to observe that there are two general ways in which an individual can contribute to an online cultural field. The fundamental idea is that on each platform there is a *primary* type of resource which a user can add to the pool, while there are also one or more *secondary* types of resources that can be contributed<sup>14</sup>. For illustration purposes, I consider a picture on Flickr and an article on Wikipedia to be primary resources, while a comment on someone else's picture on Flickr, a tag appended to a bookmark on Del.icio.us or a “this review was helpful” vote on Amazon would be secondary ones. Secondary resources are far from being irrelevant: it is through them that users weave the networks of metadata and inter-user links which underlie all user-generated content platforms. I will call "contributions" and "metacontributions" to additions made by a user to the

---

14 I am aware that the term “resource” is tricky. Berners-Lee et al (1998) define it as “the conceptual mapping to an entity or set of entities”, starting with the observation that “a resource can be anything that has identity”. I would maintain, though, that the use of the term implies a notion of “primacy” over other, “lesser” “things with identity” that could, technically, also be called resources. E.g., the concept of “the fifth and sixth paragraphs in Moby Dick” is a “conceptual mapping to [a] set of entities”. Not many, however, would consider them to be a “resource”. Hence the reliance on a distinction between “primary” and “secondary” resources in this discussion of contributions and metacontributions.

common pools of primary and secondary resources, respectively. This distinction will be used in the coming sections.

### ***1.3 A brief note on the use of “an individual's field”***

Throughout this work I will often use expressions such as "an individual's field", "her field" and "the field to which the individual belongs". To avoid confusion, it is important to note early on that I do not mean to imply that an individual can only be part of a single field at a particular point in time. In fact, an important requirement kept in mind while developing the present framework was that it should allow for individuals to simultaneously be active, important (in the context of each field to which they belong and along any of the dimensions I will introduce later) participants in several fields. Those expressions are only used for convenience and to help make the text easier to read (and write).

## **2. Literature review**

The review of relevant literature is organized into five parts. I start by covering previous work on the particular phenomenon which will be the object of my empirical work, social (or "collaborative") tagging. I then progress to work done on a related "Web 2.0" phenomenon, namely social networking websites. The next subsection provides a(n admittedly very incomplete) review – or “list of highlights” – from the literature on online communities. The fourth section provides an overall evaluation of the existing literature on the two forms of user-generated content covered in this review and explains how the present work relates to that previous research. This section concludes with a brief justification for my choice of which theoretical framework to use.

It should be noted that it was hard to locate works on user-generated content rather than on its particular forms (e.g., social bookmarking, blogs, etc). Hence the segmented format adopted for this literature review. Addressing this gap in the literature by trying to look beyond a particular form of UGC is one of the goals of the present work.

## ***2.1 Literature on social tagging***

A review of the literature on social tagging – and, most notably, on tag-based social bookmarking – platforms identified five central areas of inquiry. These are (i) the potential and limitations of social bookmarking as an information retrieval, or "resource discovery", tool; (ii) development of taxonomies of social tagging systems; (iii) motivation and the intended goals of users; (iv) individual-level vocabulary and its development; and (v) the emergence of patterns in aggregate tagging and bookmarking behavior. I will address these in order.

### **2.1.1 Social bookmarking as a resource discovery tool**

Much of the earlier scholarly attention on social bookmarking focused on discussions of its usefulness as a resource discovery tool. Mathes (2004) introduced the key comparison between formal classification systems and those based on "folksonomies". Folksonomy is the term coined by Vander Wal (2007) to refer to the emergent, distributed classification which develops through the personal classification of resources by a large group of individuals who are not constrained to the use of a controlled vocabulary. It is from the contrast between the use of a controlled, expert-defined vocabulary and the use of free keywords (which is the practice at the root of a folksonomy) that the debate on the relative merits of tag-based social bookmarking as a resource discovery tool emerged.

Mathes (2004) work alluded to the two fundamental difficulties arising from polysemy and lack of synonym control. In collaborative tagging systems, users will often employ the same term with multiple meanings and several different terms to refer to the same concept. These are two facets of the well-known "vocabulary problem" (Furnas et al 1987), although explicit recognition of its applicability to social tagging platforms would first be made by (Marlow and colleagues 2006). To this the additional problem arising from basic level variation – the fact that different individuals think of an entity at different levels of abstraction (e.g., is my neighbor's cat a "Maine coon cat" or a "feline"?) – would be added by Huber and Goldman (2004). Basic level variation poses a problem to collaborative tagging platforms but not to traditional hierarchical classification schemes. To this list Guy and Tonkin (2006) added the commonsense issue of "sloppy tags" – i.e., the use of inaccurate, misspelled or otherwise "corrupted" (e.g., by disrespecting the naming conventions followed on that platform) tags – whose usage they found to be common on both Delicious as well as Flickr.

The key advantage of social bookmarking platforms consists of achieving admittedly imperfect<sup>15</sup> classification of vast collection of resources at low cost by enabling what Tim O'Reilly (2003) called an "architecture of participation". Shirky (2005a; 2005b) maintains that the economic appeal of collaborative tagging (compared to the traditional approach to formal indexing and classification) is an important determinant of its recent success. In a similar vein, but from the perspective of the individual user, Mathes (2004) asserts that it is precisely the low cognitive costs and overall effortlessness of tagging which makes these systems successful. Comparing the classification produced by expert librarians employing a carefully defined controlled vocabulary and those produced by an online folksonomic project is an irrelevant intellectual exercise, Shirky argues: the former cannot cope with vast collections of ever-changing online resources since the cost of building and maintaining them

---

15 To which Bowker and Star would certainly ask "just what would be a universally perfect classification?"

would be prohibitive.

Davis (2005) replied to this argument that the increasing popularity of social tagging does not necessarily reduce total costs relative to those of a collection maintained by experts. Davis argues that, in social tagging platforms, the cost is merely *shifted* from the event of classification to the time of discovery. Since classification is a one-time event, while discovery will occur as many times as the need for that particular resource, it is not necessarily the case that social tagging would actually offer a lower-cost solution. The same argument was made by Merholz (2005).

Other criticism was presented by Peterson (2006), who expressed concern over the inability of social tagging systems to produce a "definitive" or coherent classification. By allowing contradictory classifications to coexist – "a folksonomy universe allows both true and false statements to coexist" –, Peterson regards social tagging as a solution inferior to that of a traditional classification based on Aristotelian categories. In his reply, Weinberger (2006) argues that it is precisely its ability to accommodate the differing views held by different groups that is the core strength of collaborative tagging; a similar argument was made by Shirky (2005a and 2005b).

Mathes (2004) had also argued for the value of "serendipity" introduced by collaborative tagging, since such systems – to use the expression of MacGregor and McCulloch (2006) – offer a greater number of "entry points" to relevant resources than a formal catalogue.

Still laboring on the dichotomy between the folksonomic and controlled-vocabulary approaches, MacGregor and McCulloch (2006) identified a "reaffirmation of controlled vocabularies" which has happened in parallel with the emergence of collaborative tagging. These authors urge library and information scientists to try to integrate lessons from these new systems – in particular, their ability to engage users and create a social environment around resource discovery – into their work rather than

dismiss them as an inferior alternative.

Finally, in a systems design work which is unrelated to the polemic described above, Bielenberg and Zacher (2005) propose a resource discovery – and, more specifically, collaborative filtering – system called "GROOP.US" based on the similarity of tagging behavior across individuals. The underlying idea is that the bookmarks of individuals with similar tagging patterns should prove fertile ground to recommend (novel) resources to other individuals in that same group, since similar tagging behavior should indicate similarity of interests and/or views.

## **2.1.2 Taxonomies of social tagging platforms**

Efforts to theorize recent developments in social tagging and bookmarking have mainly produced taxonomies of systems and types of tags. In this subsection I will review the taxonomies which have been proposed to date. This aspect of the existing literature is arguably a good illustration of the essentially ad-hoc nature of current academic treatment of this phenomenon.

Marlow et al (2006) offer the most complete such taxonomy. It concerns system design and attributes. The underlying idea is that characterization of technical aspects of the platform might help understand the behavior observed on that platform. The authors identify seven dimensions and associated "potential implications": (i) who is allowed to tag which resources (e.g., "only the content creator"); (ii) availability of "tagging aid" mechanisms which suggest tags to the user; (iii) whether the system employs a universal pool of tags ("set model") or preserves an association between a tag and the user(s) who applied it to a resource ("bag model")<sup>16</sup>; (iv) the type of object being tagged (e.g., textual vs. non-textual); (v) whether the contributor is the author of the material; and, finally, the possibility to

---

<sup>16</sup> This distinction was originally made by Vander Wal, who employed the terms "narrow" and "broad" to characterize them (respectively).

establish links between (vi) resources and (vii) users.

Sen and colleagues (2006) provide a similarly design-oriented taxonomy based on their efforts to build a collaborative tagging platform. They identified five key dimensions: (i) tag sharing (i.e., how public the tags entered by a user are); (ii) tag selection (i.e., which algorithm is employed by the system to select a subset of tags for display when the total set of tags applied to a resource is too large for presentation); (iii) item ownership (in common with Marlow et al); (iv) tag scope (also shared with Marlow and colleagues); and (v) apparently minor issues such whether and/or how to perform tag normalization across factors such as letter case, whitespace, punctuation, etc.

Earlier, Hammond and colleagues (2005) had characterized social tagging systems into a two-by-two matrix according to who creates the content being tagged (self vs. others) and who is the intended user of the tags (*idem*). According to their survey of tagging platforms, the most popular model is by far that of users tagging content created by others for their own retrieval purposes.

### **2.1.3 Motivation and intended goals of users**

The taxonomical efforts of researchers interested in collaborative tagging often adopt a different focus. That is the user. Some authors write of identifying and classifying the different motivations which lead individuals to become users of social tagging platforms; others instead produce taxonomies of (types of) tags according to their intended function. Though conceptually distinct, both types of project share the same fundamental goal of understanding the user's intended goals while using the system.

Golder and Huberman (2004) first identified that tags can serve very distinct purposes. They categorized them into seven boxes: (i) topical tags (what a resource "is about"); (ii) resource type tags

(the type of the resource, e.g., an newspaper article or a book); (iii) identification of resource creator/owner; (iv) "refining categories" (tags which are devoid of meaning if taken in isolation yet provide detail when used together with another, "main" tag); (v) subjective evaluations (e.g., "excellent"); (vi) self-reference (tags which describe the user's relation to the resource – e.g., "mystuff"); and, finally, (vii) task organizing tags (e.g., "to read"). Careful examination reveals that the tagging user is "absent" from the four first types of tags, while the remaining three types all express an aspect of the relation between the user and the resource.

Marlow et al (2006) present a "taxonomy" of motivations or incentives towards tagging, though it appears debatable whether it is indeed a taxonomy or a mere list. Marlow and colleagues list six motivations towards tagging a resource and suggest that the individual's choice of tags when tagging a particular resource will depend on their motivation(s) to do so. These are: (i) future retrieval; (ii) contribution and sharing; (iii) attract attention; (iv) play and competition (as in van Ahn's ESP Game); (v) self-presentation (by inscribing into the system something about one's relation with the resource – e.g., tagging a concert one has attended with "seen live"); and (vi) expression of opinions. The authors separate these into two broad categories, "organizational" and "social", depending on whether or not the primary focus of the user is to use the platform as a personal filing space or as a communicative/social one.

Adopting a different approach, Sen and colleagues (2006) begin by taking the seven categories of tags advanced by Golder and Huberman and then collapse them into three broader classes: "factual", "subjective" (those "express[ing] user opinions related to a [resource]") and "personal" (those that "have an intended audience of the tag applier themselves" (sic)). To these three types of tags they map five distinct "user tasks": (i) self-expression, (ii) organizing, (iii) learning, (iv) finding and (v) decision support (e.g., whether to access a particular resource). As one might expect, in a survey of users of an online movie recommender system the authors found that factual tags are most useful for the "learning"

and "finding" tasks, subjective tasks are especially useful for "self-expression" and personal tags are good for organization. In a sense, Sen et al provide a more refined treatment of the issue of motivation and intended user goals, since they distinguish between types of tags and user motivation, thus implicitly admitting that – though clearly related – these are two distinct dimensions.

Ames and Naaman (2007) looked specifically into the motivations for tagging photographs uploaded onto Flickr. Through the conduction of user interviews, their work builds on the earlier Marlow et al (2006) (on which Naaman is a co-author) and spawned a more refined taxonomy of motivations for tagging in that particular context. The result is a two-by-two matrix with the dimensions "sociality" ("whether the tag's intended usage is by the individual who took and uploaded the photo or by others") and "function" (whether the goal is to "facilitate later organization and retrieval" or "to communicate some additional context to viewers of the image").

Yew and colleagues (2006) looked at the use of collaborative tagging in a classroom setting. By coding interviews with students they distilled tags to two essential categories: "functional" ("labels that indicate some form of utility or function to the members of the class") and "content" tags. The authors point out that the high frequency of functional tags was the result of a strong social norm among students taking the course which assigned to each of those tags particular follow-up actions (e.g., the tag "opinionslugs" marked posts which were meant to initiate a discussion and requested comments from others). Although this distinction can be mapped to some of the broader taxonomies described earlier (e.g., Yew et al's "functional tags" seem, e.g., to be clearly in one pole of Ames and Naaman's "sociality" dimension, while "content tags" appear close to Sen and colleagues "factual tags"), it is too strongly tied to the particular context under study (that of tagging in teaching environments) to be generally useful.

Finally, Wash and Rader (2006a) conducted interviews of Del.icio.us users and found that in their

sample individuals the primary motivation for bookmarking and tagging resources on that platform was personal organization and their own later retrieval of resources. (The works by Golder and Huberman 2006 and Ames and Naaman 2007 concur.) Combined with the vocabulary issues mentioned in an earlier subsection, this poses particular difficulties to tag-based resource discovery, since users (at least those interviewed by Wash and Rader) appear uninterested in even trying to overcome issues of vocabulary to make their tags more helpful for others.

## **2.1.4 Individual-level vocabulary and its development**

Most of the work on social tagging looked in some way into the individual tagging practices of individual users and how these evolve over time.

Mathes (2004) was the first to hypothesize that individual tagging behavior might be affected by the tagging behavior of others (and, in particular, the tagging behavior of "specific groups of users who they share tag use with").

Sen and colleagues (2006) incorporated this idea into a model which additionally identifies two other influences on individual tagging behavior: (i) "personal tendency" (individual preferences and beliefs about the tags they apply, among which habit and previous "investment" in an established personal tagging vocabulary are especially important); (ii) community influence<sup>17</sup> (as described by Mathes); and (iii) the tag selection algorithm which determines which tags previously applied by other users are shown to the user during the act of tagging. Through an experiment they established that the

---

<sup>17</sup> The decision to use the term "community" in this context is Sen and colleagues'; following the definition of online community in Sproull and Arriaga (2007), I would not employ this word to describe an online group of mainly non-interacting individuals who neither share a common goal nor show concern for group and member welfare.

importance of personal tendency increases over time (i.e., as a user adds more bookmarks to the platform) and that the effect of the tag selection algorithm grows as the number of tags shown to the user increases.

Results by Marlow et al (2006) also suggest that "social" influence plays a role in the development of an individual user's tagging vocabulary. While Sen and colleagues detected an influence of the tag choices of the general user population on individual tagging behavior, these authors found that pairs of individuals who listed each other as "contacts" were more likely to employ the same tags than randomly chosen pairs of users. (Unlike Sen et al, who conducted an experiment, Marlow and colleagues are limited to reporting a relationship between these two variables and are thus unable to rule the possibility of them being endogenous.)

Golder and Huberman (2006) report that in their dataset markedly different usage profiles are identifiable. First, there is "not a strong relationship" between the number of unique tags applied by a user and the number of resources they have bookmarked. Some users quickly develop a set of tags they employ which remains basically unchanged as they keep bookmarking resources, while for others the number of unique tags they use grows steadily with the quantity of bookmarked resources. (Marlow et al 2006 report similarly diverse usage patterns in their study of Flickr.) Similarly distinct patterns can be observed in the usage of particular tags by particular users: some tags are (re-)used by an individual in a steady way as they bookmark more resources, while others lie dormant for a long time only to be rediscovered many bookmarks later. In trying to explain some larger-scale patterns which they observed (covered in the next subsection), the same authors hypothesize that imitation (of the tags applied by other users to the same resource, which Delicious makes visible to the user while tagging) might play an important role in explaining individual tagging behavior<sup>18</sup>.

---

18 Golder and Huberman also report an interesting, though unrelated, finding which can be useful for those interested in "sense mining" similar datasets: users appear to consistently enter tags in order of decreasing levels of abstraction, i.e.,

In their study employing semi-structured interviews of a small sample of Del.icio.us users, Wash and Rader (2006a) identified three "heuristics" used by individuals when choosing tags for a particular resource: (i) reuse tags which that individual has used before; (ii) adhere to (and, if necessary, create new) mental rules or definitions for specific tags; and (iii) choose terms which are expected to be useful later on as search terms. According to their respondents, a particular difficulty arises from the increasing difficulty in remembering relevant mental rules for the application of tags as the set of tags one has used earlier grows.

As mentioned earlier, Guy and Tonkin (2006) found that the use of "sloppy" tags is a widespread practice among users of social tagging platforms (in their dataset, 40% of tags on Flickr and 28% of tags on Del.icio.us are not present on the multilingual dictionary they employed; a critical reading of their results, though, would question whether that dictionary includes, e.g., technological and toponymic terms which are often used on those platforms).

### **2.1.5 Larger-scale/emergent patterns in social tagging platforms**

According to the work of Golder and Huberman (2006), the aggregate collection of tags applied to a particular resource exhibits a remarkably stable property. After being bookmarked fewer than 100 times, the proportions of different tags in the aggregate collection of tags applied to a specific resource becomes constant. In the words of the authors, "a nascent consensus seems to form, one that is not affected by the addition of further tags". They advance two explanations for this observation: the individual-level imitation of tags used by others (described in the earlier subsection) and the existence of shared knowledge and background among some (sub-)population of users who are thus inclined to

---

more general tags are applied first and then followed by more specific ones.

use the same tags to characterize a resource.

Mathes (2004) first hypothesized that the distribution of tag frequencies (across all resources being tagged on a given platform) would follow a power law, an idea later confirmed by Yew and colleagues (2006) in their study of classroom tagging<sup>19</sup>. He also introduced the debate regarding the "convergence" in a folksonomy: whether collaborative tagging systems have the potential to "converge on terms and foster consensus" or if as the user base grows no particular (and meaningful) classification will stand-out for each resource. Wash and Rader (2006a) made a relevant distinction between two possible meanings of the term "convergence" in discussions of collaborative tagging: convergence around a set of "correct" tags for a resource (Mathes employs the term in this sense, and the results by Golder and Huberman reported earlier confirm that this sort of convergence indeed occurs) and convergence towards a coherent use of the vocabulary in such a way that each tag is used only "in certain contexts or to refer to specific concepts". (No empirical study to date has addressed the question of whether convergence, in this latter sense, occurs in collaborative tagging platforms.)

Mathes highlights the potentially important role played by a "tight feedback loop" in the development of a shared vocabulary. By "tight feedback loop" this author refers to the common characteristic among social tagging platforms of immediately displaying to the user other items annotated with a tag the user has just applied to a resource. This immediate "confrontation" with other uses of the same tag (including by other individuals) arguably helps the development of some basic agreement on the meaning(s) of each tag. As Sen et al (2006) explain, the reason why this matters is that "social navigation may be more powerful in communities that share a common vocabulary". These authors also established by means of an experiment that the presence of a "tag recommendation system" (distinct from the feedback mechanism) can indeed promote convergence towards the use of

---

19 Cattuto et al (2006) also detected a power law distribution among tags *co-occurring* with a few high-frequency tags they had selected (e.g., "blog", "xml").

classes of tags (namely, "factual" tags) which are found to be most useful by other users.

Although not a pattern in strict sense, Yew et al (2006) observed the formation of strong social norms around the use of certain tags. In particular, tags they identified as "functional" (cf. earlier subsection on taxonomic efforts on the topic of this phenomenon) were efficacious because students were aware and followed the social norms which dictated a particular interpretation of those tags. Though Mathes (2004) also described similarly strong norms around very specific tags on Flickr – e.g., the tags "sometathatithurts" and "flicktion", both of which having particular meaning among hard-core Flickr aficionados –, it is not clear that the development of strong social norms on the meaning of individual tags is happening on a broader scale and among the generality of the user population of mainstream social tagging platforms.

Looking instead at the network of inter-user links (rather than at aggregate-level patterns in the usage of tags), Wash and Rader (2006a) report that three different kinds of users can be identified. Apparently inspired by the work of Kleinberg (1999), these authors found users who (i) have high out-degree and low in-degree ("sinks"), (ii) have high in-degree and low out-degree ("sources") and (iii) have high values for both in- and out-degrees ("hubs"). The frequency of these extreme types of users in their dataset is very low, though, with these three groups accounting for less than 1% of the total number of users in their dataset.

In turn, Cattuto and colleagues (2006) provide an interesting insight into patterns of tag co-occurrence. They find that multiple less frequent tags tend to be associated with one high-frequency tag which effectively acts as a super-category for those less frequent tags<sup>20</sup>. These authors also advance a mathematical model – which has as its fundamental characteristics there being a greater likelihood of

---

20 Those interested in building "meaning-mining" systems from tagging data could find it interesting to combine this finding with the tendency to apply tags in decreasing order of abstraction reported by Golder and Huberman (2006).

applying a tag which (i) is more frequent and/or (ii) has more recently been applied – which outputs tag distributions which closely mirror the ones in their dataset.

The work that comes closest to the main concept developed in the present work – that of an online cultural field held together by a shared interest and the desire to produce – and, along the way, value – field-relevant user-generated content – was conducted by Paolillo and colleagues (Paolillo and Penumathy 2007; Paolillo 2008). Paolillo and Penumathy (2007) looked at a dataset of 4,247 Del.icio.us bookmarks pointing to online video resources. Principal component analysis revealed the existence of an association between groups of users and particular tags they often applied. However, the authors were able neither to identify meaningful clusters of URLs (based on being tagged similarly and/or by the same users) nor to further explore the user-tag association using network analysis techniques. In personal communication, Paolillo attributes these difficulties to the sparsity of the (user, bookmarked URL) matrix<sup>21</sup> and the small size of their dataset.

In his other paper, Paolillo (2008) looked at the network of Youtube users<sup>22</sup> and its relation to clusters of users who often employ the same tags. Interestingly, and unlike what had happened with the earlier Del.icio.us dataset, the clustering produced interesting topical clusters reflective of relevant Youtube genres (e.g., Korean pop music videos, anime, religion). Also, Paolillo was able to establish that these topical clusters present very strong internal linkage and comparatively weaker ties to a more mixed "main core" of the network in which several genres cohabit (mainly, guitar playing and humor).

Finally, Mika (2007) proposes to model "emergent semantics" on collaborative tagging platforms by means of a tripartite network of users ("actors"), concepts and instances. Using network analysis

---

21 The probability of a random user having bookmarked some random URL is very low, thus making any (chance) co-occurrences stand out as "clusters" which, upon examination, most often are not interesting (or "easily characterizable").

22 Following both "friends" links as well as comments posted by a user on the page of a video posted by some other user.

techniques and by focusing on the co-occurrence of tags on each particular resource, he is able to identify clear topical clusters of tags in his dataset of Del.icio.us bookmarks (e.g., clusters of tags relating to "travel", "business", "sex" and "web design").

## ***2.2 Literature on social networking websites***

Social network websites allow users to create personal profiles and to establish links with the profiles of other users (often termed “friends”). Previous work on this phenomenon has overwhelmingly focused on issues of identity and impression management, with some attention also having been paid to analysis of the network structure of such sites.

### **2.2.1 Impression management and identity**

Boyd and colleagues (Boyd 2004; Donath and Boyd 2004; Boyd 2007; Boyd and Heer 2006; Donath 2008) are responsible for numerous studies on self-presentation on these platforms. At the center of their work lie two observations. First, individuals negotiate their identities on these sites in order to adapt to the different pressures exerted by their audiences (e.g., friends, family, work colleagues, etc). Second, the establishment of links to other individuals is itself perceived as a key aspect of online identity. That is, whom you count as “friends” – and, later established by Walther et al (2008), how attractive your friends are – is as important an aspect of a user's online identity as the more strictly “personal” information contained in the profile. Additionally, links to other users are perceived as signaling the veracity of the other information contained in the profile.

Stutzman (2006) reports on the kind of personal information revealed on these sites by college students

and writes of a “more subjective and holistic” disclosure of personal information than that which traditionally occurred on campus directories (which can arguably be compared to Facebook given the latter's partitioning of its network of users based on school affiliation).

## **2.2.2 Network aspects**

Golder and colleagues (2007) analyzed the network defined by the exchange of messages among Facebook users; they found clusters corresponding to schools and temporal patterns which correspond to the academic calendar. Ellison et al (2007) found a connection between more intense usage of Facebook and greater stocks of social capital. Lampe et al (2007) found a connection between users providing certain profile elements (e.g., high school, major) and the number of “friends” on these networks: the authors argue that such information makes it easy to identify other users with whom one shares common traits which in turn would lead to the establishment of links.

Kumar and colleagues (2006) mapped the network on one of these sites. They divide the network into three parts: a giant component of users who are accessible to each other by following inter-user links; a “middle region” comprising several isolated communities which are not connected to the giant component; and a set of singletons (users who are not connected to any other users). Perhaps their most interesting finding is that there are three distinct types of members: passive, inviters and “linkers” (those who are fully engaged with the site). The authors also find that the “middle region” is significantly large (e.g., approximately 30% of the user population on Flickr)<sup>23</sup> and that the “isolated communities” which constitute it are shaped like stars. Kumar and colleagues advance the explanation that such pattern results from, for every set of people who are invited to join the network, only a small

---

<sup>23</sup>Also of note is that its fraction of the total population stays constant over time.

subset of those who accept the invitation come to actively engage with the site.

Adamic et al (2003) analyzed an early Stanford social networking website and found correlations between (self-reported) personality types, academic majors and interests/preferences. As expected, links were found to be homophilic, with individuals being more likely to “befriend” people with shared interests.

Mislove and colleagues (2007) performed a large-scale analysis of network properties of Flickr, Youtube, LiveJournal and Orkut. This study confirmed expected properties of such networks (the distribution of links follows a power-law, the “small-world” hypothesis is verified and the network is scale-free), while also reporting that users' in- and out-degree are generally similar and the existence of a highly-connected core which bridges smaller, peripheral cores.

Of particular interest given the purpose of the current work, Liu and colleagues (Liu et al 2006; Liu 2008) superimposed an alternative network – which they termed “taste fabric” – on top of a collection of profiles obtained from Myspace and Friendster. By analyzing the correlations among different tastes and interests, the authors identified how close or far apart random pairs of interests were and thus identified different “taste ethoi” (profiles of tastes which tend to co-occur.)

Finally, Boyd and Ellison (2008) provide both an historical account of the evolution of social network websites as well as a broader review of this literature (covering, e.g., work on issues related to privacy and the relation between users' on- and off-line social networks).

## **2.3 Literature on online communities**

In the literature on online communities and electronic knowledge networks we find the vast majority of attention focused on understanding motivation for participation.

Among others, Butler (2001), Constant et al (1996), Faraj et al (forthcoming), Lakhani and von Hippel (2003), von Hippel and von Krogh (2003) and Wasko and Faraj (2000; 2005) analyze the mechanisms which lead to the motivation to contribute and thus the sustainability of online groups. In this stream of research, motivation and individual-level action are driven by a combination of self-interest, community orientation (either in the form of reciprocity or prosocial behavior), position within the network and broader-level structural properties of the network. Of the work which explicitly addresses the “sustainability” of online communities, one finds that Butler (2001) does so by focusing on the effects of community size and communication activity (neither of which seem likely to have significant impacts on “web 2.0” user-generated content platforms in which users are largely invisible to one another until they deliberately establish a link between them). Faraj et al (forthcoming), on the other hand, return to the study of individual motivation as the primary antecedent of sustainability<sup>24</sup>. Adopting a negative perspective, others (e.g., Nonnecke and Preece 2000, 2001; Preece et al 2003) engaged in studying the behavior of those who were *not* sufficiently motivated to participate yet frequented online communities – the so-called “lurkers”.

There is also significant overlap between work on the topic of online communities and research on computer-mediated communication. Given that interpersonal communication plays a minor role on the user-generated content platforms I am studying, I will refrain from addressing that research.

To summarize, work on online communities nearly always regards contribution behavior as a public gesture towards others which needs explanation. The fact that in “traditional” online communities (e.g.,

---

<sup>24</sup> And report results from attempts to run social network analysis software on particular online communities.

in online tech help groups such as those studied in Constant et al 1996) users reaped no direct, immediate benefit from their participation was (I suspect) what triggered this fascination with the issue of motivation in online communities. This interest would be reinforced when free/open-source software started being characterized as a “success story” and academics noticed that economically valuable output was being generated by unpaid volunteers (hence the studies on free/open source software, e.g., von Hippel and von Krogh 2003, von Krogh et al 2003).

Motivation, however, does not seem to be a pressing question when it comes to the current generation of user-generated content platforms, in which “contributions” have evident private benefits in addition to contributing towards the production of a public good (cf. Introduction).

## ***2.4 How the present work relates to previous research***

The first two of the previous subsections (on social tagging and social networking websites) demonstrated how fragmented our current understanding of the production and valuation of user-generated content (UGC) is. The streams of literature covered have until now refrained from addressing UGC in a minimally integrated manner, opting instead for phenomenon-specific characterizations of the different genres of UGC (social tagging, blogs, online social networking, etc) without consideration for the similarities across them.

There exist only two threads running through these literatures<sup>25</sup>. The first is the by now commonplace identification of "power laws" in the distributions of tag frequencies, number of links between blogs

---

<sup>25</sup> The use of the plural form is deliberate.

and popularity rankings of users on social networking platforms<sup>26</sup>. The second thread which "unites" these literatures is the relative popularity of basic network analysis methods and the attendant provision of descriptive network statistics.

Our discipline's extensive literature on online communities, on the other hand, tends to focus on a fundamental issue which is of less relevance in the context I propose to study and thus cannot provide much guidance.

The present work should be understood as an attempt at changing this state of affairs through the introduction of a generally applicable theoretical lens capable of looking at the production and valuation of UGC in an integrated manner.

## ***2.5 Why Bourdieu's concept of field***

One commonality of all the online phenomena grouped under the umbrella of Web 2.0 is that individuals do not engage in tight interpersonal and/or group dynamics. Compared to the more or less stable communities found in mailing lists and discussion fora, groups on Web 2.0 platforms generally are more fluid entities in which (i) little or no sustained interaction<sup>27</sup> occurs among members and (ii) neither a common identity nor a concern for group welfare is apparent. These characteristics of groups on Web 2.0 platforms distinguish them from online communities and thus invite the adoption of a different theoretical foundation for their study. Such a foundation which did not hinge on the notions of interaction and group formation was found in Bourdieu's concept of field.

---

26 The detection of "power law" distributions in connection with online phenomena apparently took off with Anderson's (2004) introduction of the concept of "the long tail".

27 Cf. definition in a footnote to the Introduction.

A second reason for my use of Bourdieu's concept and ideas is their markedly inductive character. As will be explained in the next section, Bourdieu offers a framework for inductive reasoning which is especially well-suited for exploratory theory development when studying new phenomena. In retrospect, I understand that deductive, theory-driven<sup>28</sup> work could certainly have produced the same conceptual output, yet appeared to be – at least when I started this project – less well-suited to the task at hand.

Third, my goal of understanding the social aspects of production and valuation mechanisms also makes Bourdieu's framework a “natural” choice. Here is an author who, on the one hand, has endured long-standing criticism by fellow social scientists for what many consider to be an excessively *economicist* view of social processes<sup>29</sup> and, on the other, devoted a large part of his work to understanding symbolic capital and the closely related notion of status. Therefore, if one's interest consists in understanding implicitly “economic” aspects of the social processes on user-generated content platforms (which *production* and *valuation* undeniably are) and how individuals are recognized by their peers as valuable contributors (a notion of status), then Bourdieu's framework seems to have the potential to offer useful insights. A variety of other theoretical perspectives (such as those employed in the study of online communities) would be more adequate if one were interested in understanding, e.g., motivation for participation and the psychological antecedents of a variety of behaviors.

Finally, I would like to point out that – in agreement with the second point in the paragraphs above – the use of Bourdieu's work in this proposal is mostly “instrumental”. I employ Bourdieu to guide my own inductive reasoning; the resulting concept of online cultural field and the proposed model can be understood, discussed and disagreed with independently of the value the reader finds in Bourdieu's

---

28 As will become clear in the next section, Bourdieu's ideas are closer to being a “framework” or “lens” than a “theory”.

29 Bourdieu and Wacquant (1992: ??) discuss this issue at length, citing multiple examples of such criticism. The coming section will also make clear what underlies the notion that Bourdieu thinks in economic terms.

ideas<sup>30</sup>.

### 3. Theory

Although Bourdieu's notion of field provides an excellent basis for understanding the "Web 2.0" phenomena which are the subject of this work, it does not lend itself to direct application to such phenomena. Therefore, in this section I will develop the concept of an "online cultural field". I will first review Bourdieu's concept and then identify where I deviate from it.

This is a good time to point out that Bourdieu himself avoided excessive attachment to earlier theories and advocated the (re-)use of theory as a practical tool to understand the particular phenomenon one is confronted with.

First, theoretical constructs should not be understood as rigid objects which require complete adherence on the part of the researcher. For Bourdieu, "it is better that [sociological] concepts be polymorphic, supple and adaptable, rather than defined, calibrated and used rigidly" (Bourdieu and Wacquant 1992: 23). Wacquant wrote that "[Bourdieu's] relation to concepts is a pragmatic one: he treats them as 'tool kits' designed to help him solve problems" (Bourdieu and Wacquant 1992: 30-1). He was averse to "theoretical work done for its own sake" [idem] or "treating concepts as intellectual totems" (Bourdieu and Wacquant 1992: 161), and believed that his notion of fields gained from the fact that "it promotes a mode of construction that has to be rethought anew every time" (Bourdieu and Wacquant 1992: 110). "Scientific theory" itself, he believed, "has more to gain by confronting new objects" than by adhering

---

30 This contrasts favorably with traditional deductive use of theory. In such cases, where theory is used to define the premises which guide the interpretation of empirical results, skepticism of the theory's claims and/or of its applicability to the phenomenon at hand contaminates the conclusions.

to previous theories (Bourdieu and Wacquant 1992:161).

Second, Bourdieu emphasized the importance of analogy when conducting research, particularly when faced with new phenomena which are "similar in their structure and functioning" (Bourdieu 1982: 40-1) to what has previously been studied. "Transferring that which has been established about a constructed object [...] to a whole series of new objects" is a central part of his sociological project [idem]. In that regard, Wacquant writes that Bourdieu's work provides a "sociological method" which has "transferring knowledge gleaned in one area of inquiry into another" as one of its goals (Bourdieu and Wacquant 1992: 5). Such was the importance assigned by Bourdieu to this approach to research – "analogical reasoning [...] is a powerful instrument of construction of the object" (Bourdieu and Wacquant 1992: 233) – that he identified as one of its strengths that it fostered an "analogical mode of reasoning" (Bourdieu and Wacquant 1992: 75).

Third, Bourdieu often pointed out that theory and empirical research are necessarily intimately intertwined, and that it is theory that should be guided and determined by the research object – not the other way around. "The ground for these [thinking tools, an expression which Bourdieu preferred to "theory"] lies in research, in the practical problems and puzzles encountered" (Bourdieu and Wacquant 1992: 160); they "emerged in the practicalities of the research enterprise" (Bourdieu and Wacquant 1992: 161). "Scientific theory", Bourdieu wrote, "is a temporary construct which takes shape for and by empirical work" (Bourdieu and Wacquant 1992: 161).

Therefore, one could say that the current work draws both from a theory developed by Bourdieu (his theory of fields) as well as from his broader ideas on the role and place of theory in the conduction of research (to legitimate the adaptation of the said theory to the new phenomenon of online fields).

### **3.1 Bourdieu's concept of field**

The concept of field is central to Bourdieu's work and (consistently with the third idea presented above) its development occurs throughout his work. Bourdieu has written, among many other topics, on the scientific field (Bourdieu 1999), the field of literary production (Bourdieu and Johnson 1993), the fields of journalism (Bourdieu 1998) and, more broadly, media (Champagne 1990) and the field of law (Bourdieu 1986b).

A field is, first and foremost, a social space which is held together (and defined) by the relations among the individuals (or institutions) who belong to it<sup>31</sup> and an "interest" – in a sense where "to be interested is to accord a given social game that what happens in it matters, that its stakes are important [...] and worth pursuing" (Bourdieu and Wacquant 1992: 116) – which is shared among them<sup>32</sup>.

Given the topic of the current work, I will start by covering a (for us) crucial characteristic of Bourdieu's notion of field: namely, that it is built on *relations* among individuals which do not entail any form of interaction or the existence of some social tie. It is worth reproducing two relevant passages here (emphasis added):

---

31 "A field consists of a set of objective, historical relations between positions anchored in certain forms of power (or capital)." (Bourdieu and Wacquant 1992: 16) "A field is a patterned system of objective forces [...], a relational configuration endowed with a specific gravity which it imposes on all the objects and agents which enter it" (Bourdieu and Wacquant 1992: 17). "At each moment, it is the state of the relations of force between players that defines the structure of the field." (Bourdieu and Wacquant 1992: 99)

32 "There are as many 'interests' as there are fields[;] each field simultaneously presupposes and generates a specific form of interest." (Bourdieu and Wacquant 1992: 117)

"To think in terms of field demands a conversion of the whole ordinary vision of the social world which fastens only to visible things: the individual [...]; the group [...]; and even *relations understood as interactions*, that is, as *intersubjective, actually activated connections*. [...] The notion of field presupposes a break with the realist representation which leads us to reduce the effect of the environment to the effect of direct action as actualized during an interaction." (Bourdieu 1982, pp41-2)

"[...] the real is the relational: what exist in the social world are relations – *not interactions between agents or intersubjective ties between individuals*, but objective relations which exist 'independently of individual consciousness and will'" (Bourdieu and Wacquant 1992: 97)<sup>33</sup>.

Returning to a general review of the concept of field, the next point that should be mentioned is that the position of an individual in the social space of a field is determined by her stock of different forms of "capital"<sup>34</sup>. That social position, in turn, will influence both her thoughts<sup>35</sup> as well as her behavior<sup>36</sup>.

33 Specifically asked about what distinguishes field-based analysis of the state from research on the same topic based on network analysis, Bourdieu alluded to "the distinction [...] between a structural relation which operates in a permanent and invisible fashion, and an effective relation actualized in and by a particular exchange. [...] A field is different from the more or less lasting networks through which it manifests itself. It is this structure that determines the possibility or the impossibility (or, to be more precise, the greater or lesser probability) of observing the establishment of linkages that express and sustain the existence of networks" (Bourdieu and Wacquant 1992: 113-4)

34 Cf. footnote for the earlier paragraph on the "relational" aspect of fields. Also, "positions [within a field] are objectively defined [...] by the [situation of their occupants] in the structure of the distribution of species of power (or capital)" (Bourdieu and Wacquant 1992: 97) "Two individuals endowed with an equivalent overall capital can differ in their position [...] in that one holds a lot of economic capital and little cultural capital while the other has little economic capital and large cultural assets." (Bourdieu and Wacquant 1992: 99)

35 "For the viewpoints of agents will vary systematically with the point they occupy in objective social space." (Bourdieu and Wacquant 1992: 11) "There exists a correspondence between social structures and mental structures [...]" (Bourdieu and Wacquant 1992: 12)

36 "The strategies of agents depend on their position in the field" (Bourdieu and Wacquant 1992: 101). Note that Bourdieu employs the term "strategy" in a sense that is neither intentionalist nor utilitarian (Bourdieu and Wacquant 1992: 25 and

Arguably the next most important characteristic of Bourdieu's concept of field is that of social struggle. "In a field, agents and institutions constantly struggle" (Bourdieu and Wacquant 1992: 192)." "[A field is] a socially structured space in which agents struggle, depending on the position they occupy in that space, either to change or to preserve its boundaries and form [...] A field is [...] a space of conflict and competition, the analogy being with a battlefield, in which participants vie to establish monopoly over the species of capital effective in it" (Bourdieu and Wacquant 1992: 17). "Thus, we have stakes which are for the most part the product of the competition between players" (Bourdieu and Wacquant 1992: 98). Thus, we have that "in each field hierarchy is continually contested" (Bourdieu and Wacquant 1992: 52).

Closely linked to the notion of struggle for the "profits" which are at stake in a particular field is the idea that the "interest" central to a particular field always revolves around the *generation of differences* between participants. "Participants in a field [...] constantly work to differentiate themselves from their closest rivals" (Bourdieu and Wacquant 1992: 100)<sup>37</sup>. If for Bourdieu an interest is a "socially constituted concern for, and desire to play, given social games" (Bourdieu and Wacquant 1992: 25), then an intrinsic property of those games is that their dynamics have as an outcome the generation of (social) differences between the players.

These social differences take the form of differences in the stocks of capital of individuals. Bourdieu's framework relies on "three plus one" main forms of capital. There are economic, cultural and social

---

128-9). "Social agents [...] are [...] bearers of capitals and, depending on their trajectory and on the position they occupy in the field by virtue of their endowment (volume and structure) in capital, they have a propensity to orient themselves either toward the preservation of the distribution of capital or toward the subversion of this distribution." (Bourdieu and Wacquant 1992: 108-9)

37 Bourdieu notes that the expression "constantly work to" should not be given "intentionalist" readings: "There is a production of difference which is in no way the product of a search for difference." (Bourdieu and Wacquant 1992: 100)

capitals. Cultural capital is defined as “a form of knowledge, an internalized code or a cognitive acquisition which equips the social agent with empathy towards, appreciation for or competence in deciphering cultural relations and cultural artifacts” (Bourdieu and Johnson 1992: 7)<sup>38</sup>. Social capital is defined as "the sum of the resources, actual or virtual, that accrue to an individual or a group by virtue of possessing a durable network of more or less institutionalized relationships of mutual acquaintance and recognition" (Bourdieu and Wacquant 1992: 119). In addition to these three, there is also a fundamentally distinct kind of capital. "Symbolic capital" is defined as "the form that one or another of these [three] species [of capital] takes when it is grasped through categories of perception that recognize its specific logic or [put differently] misrecognize the arbitrariness of its possession and accumulation" (Bourdieu and Wacquant 1992:119).

The concepts of field and capital are so interconnected that for our purposes it is worthwhile to go into Bourdieu's definition of the latter. "Capital", he writes (Bourdieu 1986: 46), "is accumulated labor (in its materialized form or its 'incorporated', embodied, form) which, when appropriated on a private, i.e., exclusive, basis by agents or groups of agents, enables them to appropriate social energy in the form of reified or living labor". In a different passage (Bourdieu 1990: 122), Bourdieu calls capital itself "the energy of social physics". Read in combination, these passages present two fundamental mechanisms of field dynamics: (i) capital is a product of human action (“capital is accumulated labor”) and (ii) accumulated capital enables agents to acquire more capital (“capital [...] enables [agents] to appropriate social energy”, the latter then being equated – in the second passage – with capital). This idea is given support by Bourdieu's assertion that "a species of capital is [...] efficacious in a given field both as a weapon and as a stake of struggle" (Bourdieu and Wacquant 1992: 98), which highlights this dual nature of capital – simultaneously a resource used in competition with other participants in one's field

---

38 Other definitions are found in Bourdieu's work, some of them broader in scope and which make this particular definition seem focused on *embodied* cultural capital. Bourdieu (1986: 47-48) also includes cultural artifacts (e.g., paintings or books) as well as academic credentials under the rubric of cultural capital. I opted for the definition present in the body of the text since it is closest to the notion of “cultural competence” which I will use in the next subsection.

as well as the reward for winning that competition.

Another key aspect of Bourdieu's concept of field is its deep historicism or, put differently, the central role of the past in explaining the state of a field in any point in time. A field is a space where agents map out "social trajectories" – i.e., in which they have clear past, present and future social positions. As the notion of trajectory suggests, the movement of agents through this social space is characterized by path-dependence. Present and past are so connected that it is impossible to understand one without looking into the other: "we cannot grasp the dynamics of a field", Bourdieu (Bourdieu and Wacquant 1992: 90) writes, "if not by a synchronic analysis of its structure and, simultaneously, we cannot grasp this structure without a historical, that is, genetic analysis of its constitution and [tensions]". Looking at a particular individual, Bourdieu again stresses the importance of her past trajectory to understand her current situation: "the strategies of a 'player' and everything that defines his 'game' are a function not only of the volume and structures of his capital [...], but also of the evolution over time of the volume and structure of his capital, that is, of his social trajectory" (Bourdieu and Wacquant 1992: 99). "Social agents are the product of history, of the history of the whole social field and of the accumulated experience of a path" (Bourdieu and Wacquant 1992: 136)<sup>39</sup>.

Two issues which I still have not addressed are those of the boundaries of a field and possible existence of barriers to entry. For Bourdieu, the boundaries of a field are simple to define, if only in theory: "a field is a space within which an effect of field is exercised, so that what happens to any object that traverses this space cannot be explained solely by the intrinsic properties of the object in question" (Bourdieu and Wacquant 1992: 100). Those who belong to a field do so by virtue of "their possessing a definite configuration of properties" (Bourdieu and Wacquant 1992: 107) which allows them to enter into relations with others in that field and thus begin to simultaneously contribute to

---

39 Revisit also the final quote on the paragraph on the link between position in a field and stocks of different forms of capital for another reference to the importance of individual trajectories.

producing and be affected by the forces which characterize the field. Finally, Bourdieu maintains that fields typically present some form of barrier to entry<sup>40</sup>.

### ***3.2 The concept of online cultural field***

In this subsection I will define the concept of online field and address where it deviates from Bourdieu's original notion of field. (These differences are summarized in Table 1.) An important part of this section will concern the concept of capital and of its main forms; I will return to that issue in the next subsection.

An online cultural field (OCF) is a group of individuals, and the IT-supported relations among them, who share a cultural affinity and who, by engaging in the production, dissemination and valuation of user-generated content, have an influence on the trajectory of other individuals in that field.

The first (and major) departure from the concept of field is that, while I draw on Bourdieu's theory to understand the dynamics in online cultural fields, I will not (at least initially) attempt to characterize individuals strictly in terms of their economic, social, cultural and symbolic capital(s). Following Bourdieu's own advice for applying his theory "empirically" (Bourdieu and Wacquant 1992: 161), I will use his general characterization of the notion of capital to understand which aspects of individuals or their behavior best capture the social dynamics in these online spaces. Later, I hope to investigate how these "inductive" notions map onto the forms of capital used by Bourdieu<sup>41</sup>.

The second difference seems to follow from the fact that individuals engage in online cultural fields in

---

40 "[Fields] are always marked by more or less institutionalized 'barriers to entry'" (Bourdieu and Wacquant 1992: 100)

41 In fact, a tentative mapping of some of these constructs into symbolic and cultural capital will be suggested later in this section.

*non-professional* manner: their participation in these online social forms occurs outside the scope of their professional lives<sup>42</sup>. Since this “non-professional” character is not part of the definition of an online cultural field, it is likely that this fact currently confounds our comparison of online cultural fields with the fields studied by Bourdieu (which appear to have been structured mostly around professions and spheres of economic activity). In particular, it is likely that any apparent discrepancies regarding the importance of “social struggle” in these two realms are at least partially explained by this (currently) common feature of online cultural fields.

However, and in spite of the concept of social struggle apparently not applying to online cultural fields, it should be evident that Bourdieu's closely related notion of “generation of differences” between agents applies to UGC platforms. A number of prominent features of these platforms – e.g., “most popular” lists for resources, a variety of reputation measures for users – are directly associated with a set of social processes leading to *differential valuation* of contributions and users. This opens the door for the generation of social differences between users (e.g., some agents will develop a larger following – or attain a higher reputation ranking – than others). An interesting possibility is that such status differences between users in a certain online cultural field might be associated to real struggle for (scarce, rivalrous) resources *outside* of that field; very possibly offline. Consider, e.g., the offer of lucrative book deals to prominent bloggers, or the recruitment of YouTube performers by major media distribution and production companies (McGrath 2006). In this work I will not pursue these links any further, yet they should be kept in mind as possibly revealing a connection between differences among agents in online cultural fields and effective struggle for resources.

---

42 Obviously, many individuals will engage with online fields related to their areas of professional expertise. Individuals do not, however, participate in them *as a part of their jobs*.

<b>Field</b>	<b>Online field</b>
Relational, not based on interaction/social ties	(likewise)
Shared interest	(likewise)
Agents characterized by stocks of economic, cultural, social and symbolic capital	Agents characterized by their active characteristics; the latter's correspondence to forms of capital is unclear
Agent's behavior/trajectory depends on stock of different forms of capital	Agent's behavior/trajectory depends on active characteristics
Social struggle and generation of differences as key characteristics of field dynamics	Generation of differences; the applicability of "social struggle" is unclear
Historicism / social trajectories	(likewise)
Barriers to entry	(likewise)

**Table 1: Properties of Bourdieu's concept of field and those of online cultural fields**

Now that I have identified the ways in which fields (as defined by Bourdieu) and online cultural fields differ, I will introduce the "capital-equivalent" constructs which characterize individuals in OCFs.

### ***3.3 Active characteristics in an OCF***

#### **3.3.1 Criteria for the selection of active characteristics**

A careful reading of Bourdieu's references to the different forms of capital allows one to infer the logic behind their selection and reapply it to a new phenomenon in which different constructs are necessary (emphasis added):

"A species of capital is *what is efficacious in a given field* [...], that which *allows its possessors to wield [...] an influence*, and thus to exist in the field under consideration instead of being a negligible quantity." (Bourdieu and Wacquant 1992: 98)

"The *forces that are active* in the field – and thus selected by the analyst as pertinent because *they produce the most relevant differences* – are those which define the specific capital." (Bourdieu and Wacquant 1992: 101)

"People are at once founded and legitimized to enter the field by their possessing a definite configuration of properties. One of the goals of research is to identify these *active*<sup>43</sup> *characteristics*, that is, these forms of capital." (Bourdieu and Wacquant 1992: 107-8)

"There is [...] a sort of hermeneutic circle: *in order to construct the field, one must identify the forms of specific capital that operate within it*, and to construct the forms of specific capital one must know the specific logic of the field." (Bourdieu and Wacquant 1992: 108)

All of these four quotes convey the same fundamental criterion for the identification of forms of capital: a construct needs to *be efficacious*, i.e., it must either grant the individual the ability to participate and wield an influence in the field or be operative in the generation of observable effects within the field.

At this point I will adopt the alternative expression used by Bourdieu in one of the quotes above – "active characteristics" – to refer to the "capital-equivalent" constructs in OCFs. This expression both makes clear the fundamental property of those constructs as well as avoids confusion arising from the

---

43 The (English) original reads "efficient", which is a poor (though, I am told, frequent) translation of the French term "efficient". To avoid confusion I chose to replace it with the adjective "active".

(re)use in such a different context of a term rich with history.

Before proceeding, it is worthwhile to revisit Bourdieu's definition of capital:

"Capital is **accumulated** labor (in its materialized form or its 'incorporated', embodied, form) which, when appropriated on a private, i.e., exclusive, basis by agents or groups of agents, enables them to appropriate social energy in the form of reified or living labor". (Bourdieu 1986: 46; emphasis added)

In addition to supporting the conclusion we have just arrived at ("capital [...] *enables* [agents]"), we find the important notion of capital being the product of *accumulation* over time. It can then be preserved, used and converted both within as well as across fields.

Thus, we have that our active characteristics in the context of OCFs must meet two criteria: (i) they must be efficacious, either by producing an effect observable in the present or by allowing the individual some (otherwise impossible) future course of action, and (ii) they must be the product of accumulation. These two properties are in clear agreement with the spirit of Bourdieu's theoretical framework which (as mentioned in the beginning of this section) is, at its core, a markedly historicist project.

### **3.3.2 Definition of active characteristics**

I will now introduce the constructs which together with the hypothesized relations presented in the next section make up the analytical framework proposed in this work.

While reading the remainder of this section, it is important to keep in mind the distinction between an **online cultural field** and a **platform**. Multiple online cultural fields exist on each platform: e.g., on Del.icio.us there are online cultural fields centered on “programming”, “design”, “photography”, etc. Nothing prevents a user from simultaneously participating in several online cultural fields, and this work recognizes this (and, more specifically, the possibility of a user's active characteristics in an online cultural field making themselves felt (“being efficacious”) in another online cultural field where the individual is involved; more on this later) by distinguishing between online cultural field and platform.

Two other important clarifications are justified. First, I will use the term “field” as an abbreviation of “online cultural field”. Second (and as mentioned at the end of the Introduction), I will often refer to “*the* field”, “*that* field” or “the user's field” without specifying to which field (from the several to which the user might belong) I am referring. The reason for this is that my unit of analysis is the **individual as participant in a particular field**<sup>44</sup>. Thus, throughout this section there always are an implicit “focal field” and an implicit “focal user”.

Yet, not all active characteristics live neatly within a field. External factors likely play a role in the workings of a particular field. Some of the active characteristics of such an individual in the context of a particular field, however, will be properties that exist at the platform-level while others are fully external to the platform itself (e.g., any active characteristic that was acquired offline is necessarily external). Constructs will be named accordingly: “platform-” constructs refer to properties of the individual as a user of the user-generated content platform (including, but not restricted to, this particular field), while “platform-external” constructs refer to aspects of the individual which are independent of her engagement with the UGC platform. Finally, one construct is also defined at the level of all spheres of activity (both on the platform as well as external to it) *other* than the user's

---

44 Or a (user, field) pair.

involvement in this particular field. I will qualify that construct as “field-external”. Figure 1 lays out these terms in relation to one another. When no qualifier is used, it is implied that the construct exists at the field level.

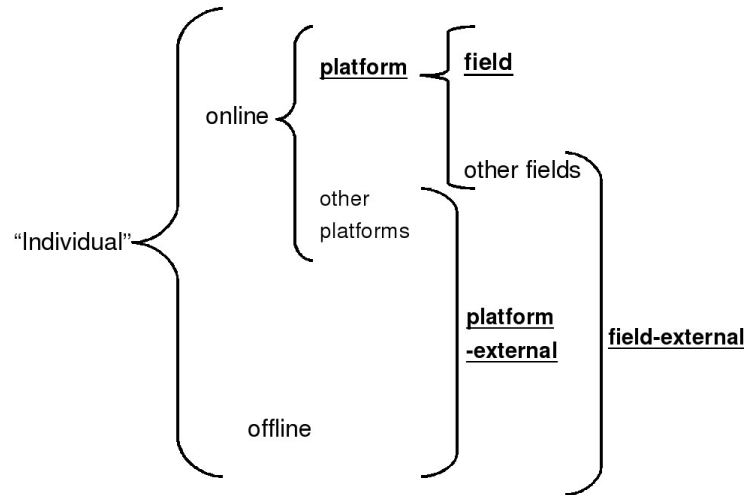


Figure 1: Construct levels

The first step was to consider the user as a *producer* of content. As a producer, an individual in a field both contributes as well as the object of evaluation (or status-conferral) by other users in the field. With this in mind, I identified the following active characteristics:

- (1) **Status:** All online social platforms support user accounts. The actions (both on that platform as well as elsewhere) of account holders will lead some accounts to become, over time, more popular and/or better-known than others. This construct captures the extent to which other users within a particular field express interest and/or appreciation for a user. It is important to note that here I am interested in *user-level* measures of a user's reputation – I will consider valuations of the user's individual contributions later. The most direct operationalization would be a measure of the “in-links” from other users in the field which a user has accumulated (as

compared to the scores assigned to a user's contributions, which would be *contribution-level* valuations).

(2) **Cultural competence:** the combination of a user's knowledge and skills which enable her to make contributions to an online cultural field. The relevant knowledge is knowledge about the topic of interest which is central to the field (e.g., programming) and on which the user can draw to make contributions (e.g., photos that the user has taken and can upload to Flickr; or knowledge of valuable articles on the web which she might place on Del.icio.us). Relevant skills include the ability to create or locate resources for contribution (by analogy, the *ability* to take photos for contribution to Flickr or the *ability* to locate interesting web content); the ability to infer which contributions would be most valued by other participants in the field; the ability to annotate content in such a way that will make it more accessible by and/or useful to other field members; and technical proficiency in the use of the platform<sup>45</sup>. For analytical purposes, I propose to break this construct down into: a unobservable, exogenous **platform-external cultural competence** accumulated outside the UGC platform<sup>46</sup> and which I will assume to be a fixed property of the individual; and an observable, endogenous (to the platform) **platform-developed cultural competence** accumulated through engagement in both the field under study as well as other fields on this particular platform<sup>47</sup>.

(3) **Field-external identity characteristics:** this construct encapsulates all aspects of a user's life

---

45 With the exception of the first of these skills (“ability to create or locate resources for contribution”), one can think of the acquisition of these skills as a sort of “meta-learning”: by developing them the user is basically learning the “rules of the field” which enable her to be a more valued member without significantly exerting more effort.

46 An example would be the ability to make relevant contributions to the field of “programming” on Del.icio.us due to one's long career in software development.

47 An example would be the additional ability of a member of the “programming” field on Del.icio.us to make relevant contributions to that field by virtue of her engagement both in that field as well as in the fields of “design” and “mechanical engineering” (also on that platform). The distinction between these two sources of platform-developed cultural competence will become relevant when operationalizing this construct.

outside of the online cultural field which are not directly related to her ability to make contributions to the field. It can be broken down into: **status accumulated in other fields** (on the same platform) and **platform-external identity characteristics**. The former can translate into status within a particular field because users in that field might recognize the individual as a valued contributor in some other (possibly related) field which they, at some point, frequented, and thus confer her greater status within the focal field. The other component, platform-external identity characteristics, encompasses all aspects of an individual's life outside of the platform which are not directly related to her ability to make contributions to the online cultural field. Examples would be one's demographics, interests and preferences, professional background and status accumulated on any other sphere of life. All such identity characteristics are likely to be easily convertible into status within a cultural field. E.g., Bill Gates would most likely accumulate significant status in an online cultural field centered on IT solely for the reason of him being the famous founder of Microsoft. Or, as Foreman et al (2006) showed, merely revealing (a not famous) name and location on such a platform is rewarded by others with relatively more positive evaluations. As different as these two cases are, they both illustrate how aspects external to one's participation in a field might be efficacious within the field

Since at the center of online fields lie not interpersonal communication and interaction but instead the creation, exchange and valuation of user-generated content, it is also important to look into the properties of a user's contributions and metacontributions. In particular, and in line with the theoretical work in the previous section, I am interested in identifying aspects of a user's contributions and metacontributions which will accumulate over time and which can be efficacious in the field (e.g., by making the user's contributions more valuable to others and possibly conferring the user higher status). As mentioned in the Introduction, what constitutes a contribution and metacontribution is platform-specific (e.g., on Flickr uploading a picture is a "contribution" while appending a comment to a photo

taken by someone else is a “metacontribution”). With this in mind, I identified the following aspects of contributions and metacontributions which constitute active characteristics:

- (4) **Quantity:** total number of contributions or metacontributions made by this user in the field
  
- (5) **Originality:** it is often possible (and, on many platforms, likely) for multiple individuals to contribute duplicates of the same resource. Similarly, metacontributions can be original (e.g., the first comment to make a particular observation) or not. This construct is intended to reflect the extent to which an individual contributions and metacontributions add to the field through their originality.
  
- (6) **Uniqueness:** while the originality construct captures the notion of an individual being the *first* to make a particular contribution or metacontribution, uniqueness evaluates an individual's contributions or metacontributions to the field by asking "to which extent is this contribution (or metacontribution) something that would otherwise would not exist?". Uniqueness arguably offers a stronger claim to being a valuable participant than originality. (Compare "if it hadn't been for me, X wouldn't even be available" with "I was the first to contribute X".)<sup>48</sup>

It should be noted that both originality as well as uniqueness are only meaningful on platforms which practice little or no control over duplicates (at the time of writing, this is the case on all major platforms mentioned in the previous section). Furthermore, the uniqueness construct also requires a platform on which there is a distinction between user-level pools of resources (e.g., "joe213's bookmarks") and the common pool of resources (e.g., "all bookmarks").

---

48 An interesting exception might occur in online fields focusing on the creation and exchange of very time-sensitive resources (e.g., user-contributed news sites such as Digg). There, being the first to contribute a particular resource would reflect positively on the individual, while the question of whether or not someone else contributed one or more duplicates at a later date – i.e., whether or not it was a unique contribution – would be irrelevant when discussing the worth of the original contribution.

(7) **Prototypicality**: measures the extent to which the contribution or metacontribution is aligned with the prevailing interests, preferences and/or view(s) of the field. Being more prototypical (e.g., contributing a bookmark for an article on web programming to the OCF on Del.icio.us focusing on the Python programming language, which (hypothetically) exhibits an established interest in the use of that programming language for web development, or a picture of a scene in the historical downtown of Lisbon to the Flickr OCF on “pictures of Portugal”, which would have shown a similar inclination towards photographs of those neighborhoods) might make an individual either more valued (by addressing a preference of the field) or less so (perhaps the field values more off-stream contributions). It should be noted that I have not fully “fleshed out” the relation between prototypicality and the two earlier constructs (originality and uniqueness). My current idea is that originality and uniqueness are more strictly oriented towards the “content” of the contribution (adopting the two earlier examples: “has this article on Python web programming been bookmarked by other individuals in the field?” and “are there other shots of this particular street view already available?”), while prototypicality instead addresses the relation between a contribution (or metacontribution) and a broader notion of the field's interests and preferences.

(8) **Revealed interestingness**: this construct captures the evaluations of this contribution<sup>49</sup> which have been expressed by other users. Such expressions of valuation are made through the metacontributions of users (e.g., “helpful” votes on Amazon, or adding a picture as a “favorite” on Flickr).

Finally, one must not forget the fundamental duality which underlies user-generated content platforms:

---

<sup>49</sup> No platform that I know of, with the exception of Slashdot (which has a “meta-moderation” system), allows users to rate the metacontributions of others. Hence, I refrain from extending this construct to cover metacontributions.

users are simultaneously producers *and* consumers of content. With this in mind, I defined the following consumption-side active characteristic:

- (9) **Content viewed:** a measure of the resources contributed by others which the user has accessed while using the platform. Importantly, I will distinguish between **content viewed within the field** and **content viewed in other fields**.

This completes the list of active characteristics which I have, until now, identified. I regard it as provisional and look forward to revising it.

## 4. Proposed model

In this section I will propose a dynamic model of how the active characteristics identified in the earlier section might feed into each other and thus sustain the social dynamics within fields on user-generated content platforms.

The model is depicted in Figure 2. It is worth once again reminding the reader that the unit of analysis is the individual as participant in a particular online cultural field. Additionally, when an arrow begins and/or ends in an outer box (i.e., one containing other boxes) that should be understood as an abbreviation: that arrow represents multiple arrows beginning and/or ending at the individual inner boxes contained in the outer box.

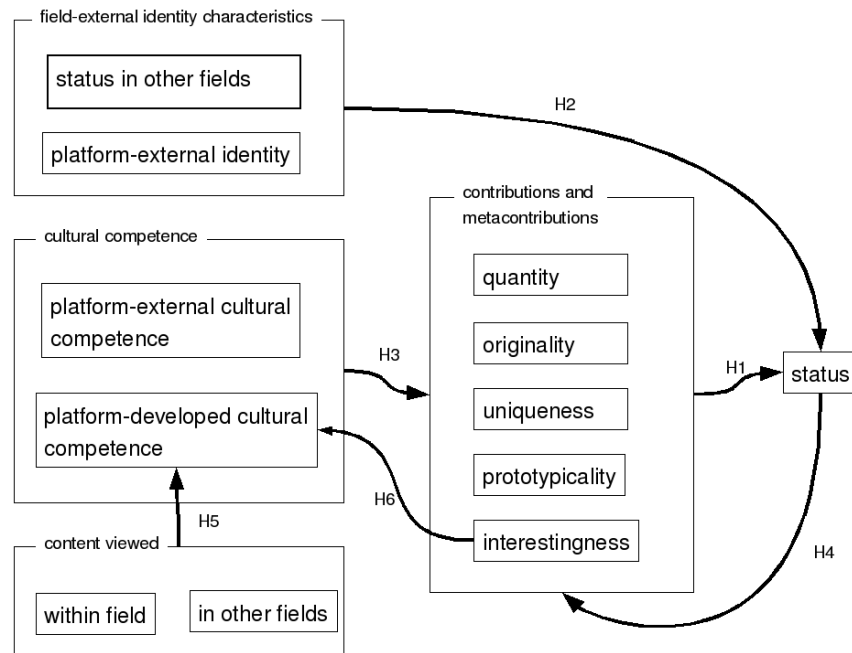


Figure 2: Proposed model

Broadly, I propose to investigate six main types of links between constructs:

**(H1) Properties of a user's contributions and metacontributions will confer status to that user.**

**(H2) “Capital” other than cultural competence and accumulated outside the field<sup>50</sup> will also grant the user status within the field.**

**(H3) Cultural competence will enable the user to make a stream of contributions and metacontributions with the properties listed earlier.**

<sup>50</sup> In both the previous subsection as well as in the model depiction in Figure 3 this aggregate form of capital is termed “field-external identity characteristics”.

**(H4) Status will induce a user to further make a stream of contributions with those same properties.** The underlying idea is that, possibly driven by a desire to preserve the earlier-acquired status, a user will feel motivated to continue making (more) contributions.

**(H5) Consumption of content contributed by others will contribute to the development of cultural competence.** The underlying idea here is that, through the consumption of other resources on the site (browsing), a user acquires knowledge and skills which will help her, at a later stage, make (more) valuable contributions. As mentioned in a footnote when describing this construct in the previous subsection, one can conceive of a user acquiring competence along two dimensions through the consumption of the contributions others have made: “material” learning about the topic which is the focus of the field and “meta” learning of what the general interests and preferences in the field are. (I will not try to decouple the two.)

**(H6) Observation of the interestingness scores assigned to one's contribution stream will help develop cultural competence.** This channel of learning will only produce cultural competence of the “meta” variety: observing a stream of feedback for individual contributions will allow a user to induce what is perceived as valuable in the field.

Both active characteristics characterized as “platform-external” (both on the left-hand side of Figure 2) are assumed, when studying a given individual in a particular field, to be fixed. That evidently is a simplification. An individual's “platform-external identity characteristics” and “platform-external cultural competence” are liable to evolve over time: e.g., a person might move a different job and/or might acquire new cultural competences offline. Since neither of these constructs is observable by me and any changes to them are likely to only make themselves felt in the online cultural field over large(r) periods of time, I will regard them as exogenous and fixed for any given individual within a particular

field.

## 5. Empirical work to be done and methods

This section will outline the steps involved in conducting an empirical analysis of online cultural fields on a user-generated content platform. I plan to conduct empirical work on two different datasets, but only have data and a clear(er) plan for one of them. At this point, I intend to consider solely contributions (not metacontributions), although the ideas outlined here should be extensible to permit the inclusion of properties of metacontributions into the model.

### 5.1 Choice of sites and data collection

The first dataset was obtained by crawling the social bookmarking platform Del.icio.us during November and the first weeks of December of 2007. My decision to gather data from this site was mainly based on what I identified in the Introduction as the initial reason for my interest in online cultural fields: namely, that they can be very much *unlike* online communities. Del.icio.us offers no space for interaction (not even the uni-directional posting of comments on others' contributions) and users appear, on first sight<sup>51</sup>, to be behaving in strictly individual manner.

I have assembled a collection of approximately 550,000 user profiles containing the bookmarks, tags and network connections (both in- as well as out-bound) for those users. The crawl was seeded by obtaining the list of 120 “most popular” tags listed on the site at the beginning of the crawl; from there, I expanded the root set of tags to include all tags identified by the platform as “related” to the seed tags.

---

<sup>51</sup> But not in fact, as earlier work on social bookmarking platforms has shown. Cf. the sub-subsection on “emergent aspects” on collaborative tagging platforms in the literature review.

Then I started crawling all profiles of users which had used any of the tags in that expanded root set, while also crawling the profiles of all users who had a network link to a user in that earlier set. I have yet to extract features from the HTML and cannot, for that reason, present any descriptive statistics at this point.

Since performing the crawl, I have become aware of the desirability of having longitudinal data. The data collected contains creation timestamps for all bookmarks, yet the HTML does not contain that important piece of information for links between users. As a result, I do not know when in- out-bound links were created between users. I intend to begin crawling Del.icio.us regularly in order to observe the appearance (and disappearance) of links between users and thus be able to infer a(n approximate) date for their creation (or removal). An alternative idea which has recently crossed my mind is to use the data I already have by making two discussable assumptions: that (i) a user establishing a link to another user decides to do so upon finding several interesting bookmarks in the latter's profile; and that (ii) at the time she creates the link, the subscribing user will also add several of the other user's bookmarks to her own profile. Having timestamps for the creation of bookmarks by both users and working under this assumption, I could then infer that an "A->B" link had been established when a certain amount of "bookmark importing" in that direction was observable over a short period of time. I would appreciate feedback on whether this seems an acceptable idea or if I should instead stick to the plan of crawling Del.icio.us anew (and doing so regularly).

The second dataset has yet to be collected. The crawler which I wrote should be easily modifiable to do so. I am currently thinking of using Flickr; I would appreciate any ideas or suggestions.

## ***5.2 Identification of a field, its participants and its contributions***

The first step in the analysis consists in identifying the shared interest(s) which lie at the center of each field. In the case of Del.icio.us, I will do so by identifying clusters of semantically related tags which are often applied to the same resources. Each of these “semantic” clusters of tags will be interpreted as corresponding to a shared interest on that platform. Such a step is required since (as pointed out in a footnote to the Introduction) Del.icio.us does not offer a feature that “explicitly” supports online cultural fields.

Second, for the study of each online cultural field I will isolate all users on the platform who are frequent participants. I will most likely set a minimum threshold of engagement with the field (in terms, e.g., of number of contributions to the field) to exclude from consideration individuals who are not active participants. On Del.icio.us, a user will be considered (or not) as a member of a field depending on the overlap between the set of tags which she most frequently uses and the tags which characterize the field<sup>52</sup>.

Third, there is an issue of defining to which field(s) an individual's contributions belong. (This problem evidently only poses itself when an individual belongs to multiple fields.) We need to assign each contribution to a field in order to decide whether or not to take it into account when computing the aggregate properties of the user's stream of contributions in a given field (those listed in the box “contributions and metacontributions” in Figure 2). I propose to do such an assignment based on whether or not the user tagged the resource (bookmark) with a tag characteristic of that online cultural field (see first paragraph of this subsection).

---

<sup>52</sup> This corresponds to the proposed operationalization of the prototypicality of a user's contributions within a field.

### 5.3 Operationalization of constructs

I will then compute measures for the active characteristics developed in the previous section for each (individual, field) pair. It should be emphasized again that individuals can simultaneously belong to multiple fields, and that most of the active characteristics will take on different values *for the same individual* as I analyze different fields. The reason for this is that the constructs are intrinsically relational and field-specific, and thus depend on the definition of which resources and individuals belong to the field currently under analysis.

The proposed operationalizations of constructs in the Del.icio.us dataset are presented in Figure 3. Underlying ideas and foreseen issues:

- (1) **Inter-user links** take on two interpretations: an **in-link** confers the focal user status, showing that another user (either in the focal field or outside of it) finds her activities worth following; an **out-link** is used in the Del.icio.us dataset as proxy for consumption/viewing of bookmarks contributed by others. In particular, each out-link to fellow members of the same field is interpreted as indicating consumption of bookmarks within the field; out-links which cross field boundaries are seen as indicative of the consumption of bookmarks outside that field<sup>53</sup>.

---

53 A problem with this idea is that inter-user links between two individuals who are present together in at least one field become difficult to classify: are these links across a particular field's boundaries or not? I propose to treat them as if they were links *within a field*, since some similarity of interests between the two individuals has already been evidenced by virtue of them being collocated in at least one field.

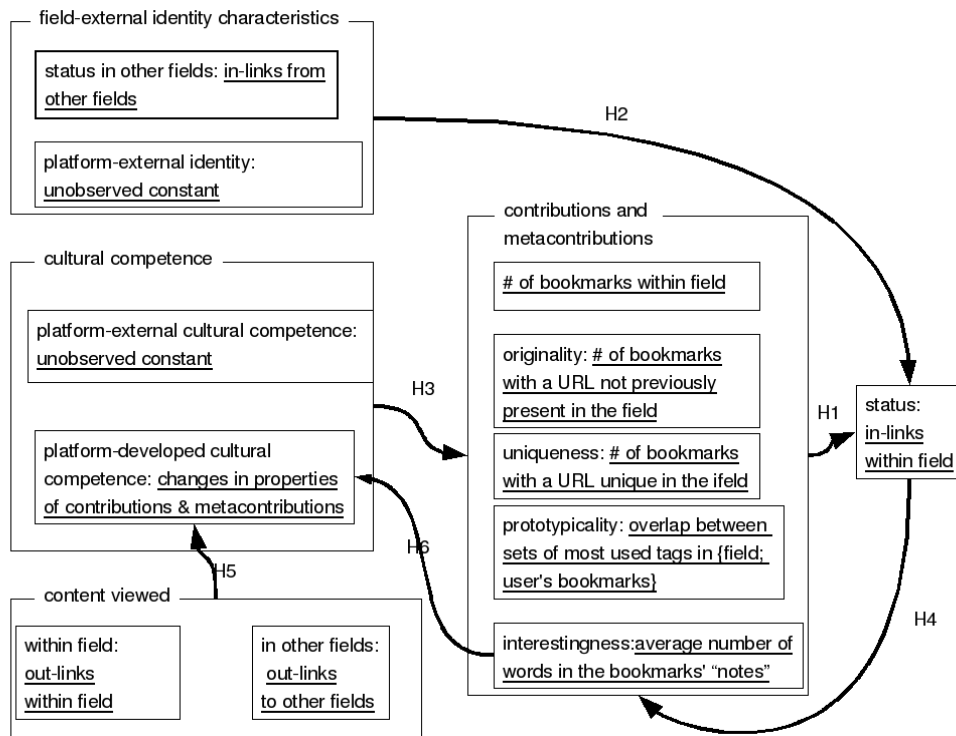


Figure 3: Proposed operationalization on Del.icio.us

- (2) **Interestingness:** Del.icio.us lacks support for users to express interest in a particular bookmark contributed by others (other than by adding it to their own stream of bookmarks). It does, however, allow a user to add notes to each bookmark. I propose to approximate the other users' “interestingness” rating of a user's stream of bookmarks with the average number of words employed by other field-participants in their notes for that URL. An alternative<sup>54</sup> idea is to consider the number of other individuals who have bookmarked that same URL – thus making “interestingness” the inverse of “uniqueness”.
- (3) **Platform-developed cultural competence:** since this is not directly observable, yet in the model it is the only non-static determinant of contribution properties (the other being the fixed

54 Or complementary, if one combines them into a single score.

platform-external cultural competence), I propose to measure developments in the cultural competence of a user due to her use of the platform by computing a composite score of the change in the properties of her contributions. The idea is that if the user is developing a cultural competence that makes her a better contributor (and platform use is the only source for such learning in this model), then it should be possible to capture the development of that competence by measuring the evolution in the observable properties of her contributions.

#### ***5.4 Testing relations between constructs***

In the absence of longitudinal data, I will temporarily investigate correlations between the constructs in the model.

### **6. Expected contributions**

With this work I hope to contribute to our understanding of how user-generated content (UGC) is produced and valued by online groups in several ways.

First, this work proposes a dynamic and relatively comprehensive model of social processes in these online spaces. Earlier research on UGC platforms has, for the most part, looked at particular pairs of constructs in isolation and refrained from theorizing the process(es) through which groups on these platforms thrive and sustain themselves. The literature on online communities, on the other hand, has overwhelmingly focused on issues of motivation for participation and interpersonal communication. This work proposes to address this gap.

Second, what I propose constitutes one of the first adaptations of a theoretical framework of

demonstrated usefulness in the social sciences (Bourdieu's concept of field) to the online realm. This work shows how Bourdieu's constructs can be translated into their online equivalents and how field dynamics can be applied to help explain social phenomena on the Internet.

Both of these factors should enable this work to be generative of new research. From a practical perspective, I would highlight three additional contributions.

First, the model advanced here can help develop thriving user-generated content platforms. It identifies different contribution patterns and their respective valuations expressed by others on the platform. This information should enable those who operate the platform to, e.g., incentivize the contribution patterns which are deemed the most helpful by rewarding the users who correspond to them. Similarly, the same information could be used to improve system design in such a way that would promote valuable contribution behavior. Finally, this model should enable the development of innovative content recommendation systems on these platforms. Having as a basis a combination of rich multidimensional constructs for both user as well as contribution characteristics, developers should be able to provide users with more interesting results than those generated through the use of aggregate-level statistics.

Second, the model advanced in this work offers electronic retailers a number of possibilities to improve their e-commerce strategies. Research on user-contributed reviews or, more broadly, “word-of-mouth” has established a link between user-contributed resources on products (e.g., Godes and Mayzlin 2004 on review properties; Duan et al (2006) on the volume of reviews) and sales volume. By determining which kinds of contributions and properties of their author(s) are tied to higher sales, electronic retailers stand to gain by, e.g., conferring greater prominence to such contributions. They could also benefit from designing their sites in such a way that they would be sure to support all those “active characteristics” which induce greater sales. Perhaps even more importantly, this framework could guide the site operators towards exploring the potential of thinking of their user population as being divided

between different fields of interests. E.g., an individual who is an important contributor of content on home appliances might not enjoy deserve similar status among consumers interested in jazz music. This multidimensionality of a user's presence on user-generated content platforms remains, to the best of my knowledge, unexplored.

Third, this model can be used as a basis for innovative pricing strategies for online advertising. By developing a metric of the value of a particular contribution (based on a combination of user and contribution characteristics), platform operators adjust the price of advertising space in accordance with the estimated value of that particular piece of content. Such an approach might compare favorably with the commonly used “pay-per-click”, since it would not be merely *reactive* to the (relative) popularity of a resource.