

Explaining and Capturing CoP Value

Miguel CORNEJO CASTRO¹

¹Macuarium Network SL, C/ Sandoval 25, 4^o izda, 28010 Madrid, Spain
Tel: +34 91593293; Email: miguel@macuarium.com

Abstract. This paper will attempt to explain the Macuarium Quadrant methodology for managing CoP-derived value, and advance some ideas on means to capture and trade it. (If approved for KMSS presentation, this paper should spawn an open conversation and debate on the subject).

1. Introduction

In previous papers (<http://www.macuarium.com/econoclasta/network>) we have discussed a novel way of mapping the value generated by a CoP to its different stakeholders. This will be reviewed and explained in the next chapters. We will then examine the practical implications of this for community management, and for a most practical aspect of it: who to charge for what, and how to make sure that they will pay. The paper is based in extensive experience at the innovative Macuarium.com CoP system.

2. Knowledge and value

Knowledge Management research is broadly divisible amongst two schools: those who take into account the economics of it, and those who don't.

The first is most commonly represented by "action-researchers", practitioners who have real, hands-on experience of trying to implement methodologies to improve the management of knowledge. These people will often try to link what they do to the criteria by which it is judged. Those criteria are ultimately determined by the goals of the entity that sponsors their activity. And those goals, more often than not, are mostly economic.

The economic goals of typical entities are quite summarizable in two families: to create as much value as possible for their customers, and then capture it by making them pay the highest possible prices; and to communicate to the financial markets an image of themselves that contains every ounce of positive information, in order to make their share price rise. In other words: how to generate earnings, and how to translate potential earnings into share price [1].

Thus, we could say that studying Knowledge Management from a value-creation perspective is a sort of "teleological" choice: we study the subject in the light of what we expect to achieve from the subject.

Many other avenues of study are possible, and many intersect with this. But we believe that many practical consequences may be drawn from our destination. And we believe practical consequences to be a worthy reason for study: we define ourselves as practitioners.

So we will now try to link Knowledge Management using Communities of Practice to the first family of economic goals: how CoPs accrue value for users, how to map it, and how to capture that value in a way that the organization can measure... and eventually charge for.

2.1 A working definition of Knowledge

As stated in previous white papers (<http://www.macuarium.com/econoclasta/network>) we will be using a working definition of knowledge that may or may not fit the bill of every researcher: it certainly does not address the four Greek perspectives on the subject [2]. But, in our teleological pursuit, it should suffice.

As a working model, we will define “knowledge” as the set of information that enables the receiver to satisfy a need. It is the sum of data and context necessary for a person to perform a desired activity (physical, commercial, mental or emotional). Therefore it is knowledge only when it is actionable. It is knowledge when it enables. This enablement creates a degree of satisfaction also called “utility”.

If it is not good for anything but the person can still understand it, we will just call it “information”. If the person can’t even make sense of it, it will be merely considered “data”.

2.2 Utility and value

We have used the concept of “utility” in its widest acception. We have made a hypothesis about knowledge and utility: that knowledge is inherently useful.

The last step is the link between usefulness or utility, and value. We believe that “value” is a quality of things that is context- and subject- sensitive: it is a measure of the thing’s usefulness in a given situation for a given person. The expression of this measure can be made through other goods: a certain quantity of one good will be as appreciated by the subject as a certain quantity of the other. An entity can find its particular equivalence for cannons and plows. Or for, say, knowledge and money.

This means that the same things may be of different value to two different people at the same time. This enables trade: both will want to maximize their total utility, and will exchange goods until they have achieved their respective maximums [3].

When one of the goods exchanged is money, we talk of the amount exchanged as “price” [4]. Money plays the role of universal measurement good: most frequently we can express value as an amount of it.

2.3 Capturing value and the requisites of trade

When we use some goods or services with a given cost to produce something, which in turn is valued by a prospective buyer at a higher price than the sum of those costs, we say that we have “created value”. When the trade takes place and we get the money equivalent of that value, we say that we have “captured” that value.

This looks simple, but isn’t. Let’s approximate some of the difficulties.

“Something tradable” is a working synonym for “an asset”. For something to be an asset, it will need to be:

- Quantifiable. At least at an aggregate level, we need the ability to measure it in order to use it, parcel it and trade it. When something can’t be quantified it is eminently non-tradable.
- Valuable. If we mean to trade it, we must be able to build an expression of its value in a common exchange good. That is to say, we need to be able to give it a value independent from its cost.
- Tradable. An asset needs to be able to change owners, and to do so without losing its value-creating qualities. If the good is only valuable for one entity, its quality as an asset is quite restricted.

- Private. The underlying, necessary condition for trade is that the good must be subject to the exclusive and recognized right of exploitation and trade by a particular entity [5], be it a person, a corporation, a state, or any other.

Not all goods comply with all four conditions. This doesn't necessarily eliminate them as assets, but places some restrictions on the way in which they behave. We will see some in a short while, as they affect knowledge in no uncertain manner.

2.4 Knowledge as an asset (the myth, the topic, the fact)

It is a frequently repeated assertion that knowledge is a "different" type of asset: sharing it, it is said, does not deprive the original owner, while it enriches the second.

This is evidently false. There is one quality of this asset that very frequently changes when it is thus shared: its value. As we have said, knowledge is an enabler: it enables the holder to do something productive (allows to create value). The value of knowledge has a scarcity component and will often be lower as more people share it, through an increase in the available offer of the resulting good.

The value of knowledge also frequently has a competitive component, since it enables the owner to do things in a different way from those who don't have it. Differentiation can be a key goal in a competitive market environment, and thus this component may have a very serious weight. If so, sharing the knowledge will very seriously depreciate it.

Once the "different type of asset" assertion is stripped as a myth, some topics become unsustainable. Most of the explanations for the difficulty in inducing knowledge sharing deal solely with trust-building and social factors. This is radically incomplete: sharing unilaterally is just a misleading name for "giving away", and most rational humans and entities will want to trade, not give, their valuable asset.

Most types of knowledge have value and behave as a fairly normal type of asset. As assets, they are tradable. In appropriate situations, people will be ready to trade knowledge for other goods in a combination that is meaningful for them.

3. Knowledge Communities

It would be quite interesting to try and apply these hypotheses to other KM tools and methodologies, but for the time being we have been focusing on the one in which more different agents interact and more different goals are pursued.

We define "knowledge communities" as conceptual gatherings of people with a shared interest in a subject, who exchange information about it in order to learn and to solve problems. At this definition level, we don't attach importance to the place or means of interaction, nor to the quality of the community's subject matter.

These Communities will come in all shapes and sizes, from amorphous debate forums lost in the biggest Internet portals to tightly-structured project development environments within a corporation, to all sorts of mailing lists and blog networks. The stakeholders will vary just as widely, spanning all parties interested in the Communities working: not just users and owners, but the external entities affected by their results and by their costs.

The only really defining requirement is that they effect that phenomenon which we were analysing above: their members exchange knowledge.

We will not go now into the mechanics of transmission of tacit knowledge and expertise, or the role of Communities in innovation: they are enough for several books. It must suffice us by now that these Communities are working knowledge exchanges.

2.1 Shapes, tools and the mystery of investment

In our experience, the character and effectiveness of a Community stems from its vision and management: the rules it works with, the goals it works toward, and the mechanisms through which it steers its course. All other factors are, in a metaphysical sense, “accidental”: even in Communities that are closely based in a particular medium (a coffeehouse, a chat, or a bulletin board), the group of people could collaborate just as well on a different platform if it fits their values, rules and ends. Even the language divide has proven to be surmountable.

Tools will be chosen and used to enhance the effectiveness of the Community. Those tools can range from shoestring operations to collaboration megasites, and can enable collaboration of different kinds and varying effectiveness. Usually, the function will develop the form it needs; we have yet to see a tool that defines a Community.

But tools and shapes have associated costs (in money, time, or other resources). Thus, since tools are a result of function, we can conclude that a useful tool is that whose cost is less than the benefit it provides.

And this leads the way to a curious question: how can a Community channel resources? How can it measure the value created by a particular tool, a particular service, a particular content? How can it raise the resources to acquire it?

How can it capture the value it creates?

2.2 Going objective

In many cases, at present, the convenience of such tools is judged from on high: the Community will often be operating as part of an entity’s intranet, which will accord it resources almost regardless of needs or profitability (after all, how can they be known?) based on assumptions, opinions or indicators.

Frequently, also, the Community will be an owner-less affair constituted from independent individuals linked by the web. There, the measurability of value is an even more distant dream... and its capture is most often undreamt of.

Some Communities, if few, attempt to capture the value they create. The original The Well community charges for membership. Some mailing-list based communities do, as well. Many companies and institutions participate in “clubs” or “poles” which charge their members for the right to take part in different sorts of knowledge exchanges. In most cases, this charge is cost-related, not value-related, and hardly efficient.

What we will attempt is first a classification of value generated for different stakeholders by different types of Community, separated according to their vision and management: goals and procedures, authority and stakeholders.

Then we will analyse this value to see how it may be captured.

4. The Macuarium Value Quadrant

On previous papers (see <http://www.macuarium.com/econoclasta/network>) we presented a tool that has been of use in the development and consulting projects of Macuarium Network. It was developed in 2002 by Miguel Cornejo, with assistance from Manuel Medina, and remains a valued piece of the company’s intellectual property and a pillar of its services offer.

This Quadrant attempts to map value creation to users and stakeholders, and thus serve as guide for the capture of it when possible. It is specially useful when assessing ebusiness initiatives and business plans, and may give us an insight on the subject of Knowledge Communities.

3.1 Value for organizations and for individuals

We won't be modeling organizations in any complexity. But we will assume that they make decisions based on their own operational needs: maximizing utility by furthering their own goals, such as benefit and growth. As far as we are interested today, entities improve towards those goals through the enablement of their people. This is what fosters innovation and guarantees quality in their processes.

This enablement is achieved through the acquisition of the relevant types of knowledge and expertise. Communities can play a key role in their transmission (and even frequently in their creation).

As we have mentioned above, individual both inside and outside an organization can derive value for themselves through the acquisition of knowledge at a cost lower than the utility they can extract from it.

Organizations and individuals react to knowledge-derived value in two different ways, which enables us to classify the value in two distinct categories.

3.2 Objective value / direct value

We observe a type of behaviour when the value of the knowledge acquired by the individual is comparable with that derived from other productive factors used in the development of the person's activity. In other words, when the knowledge received can be directly applied to the improvement of the person's economic situation or even in the execution of the person's job... or when the knowledge's quality as an asset is high.

This area will usually cover such knowledge as methodologies, precedents, solutions to problems, professional growth tools, tradable expertise. We call this type of value "Objective value", since it is measurable, quantifiable, and individuals can rationally act upon it.

The equivalent type of knowledge for organizations will be that which the organization can perceive and measure, and put in direct relationship with improvements in its processes and operations. It will be knowledge that can be traceable to results: in other words, it will have high asset quality.

We call this type of value "Direct value", since it can be ascribed to a particular body of knowledge in the same way than traditional cost accounting traces direct costs to a particular center.

3.3 Subjective value / indirect value

Individuals and organizations don't react to all knowledge in the same way, and not just by a matter of degree. The Quadrant works on the assumption of two different types of value generating two different types of attitude.

The second type, or "subjective value", is provided by knowledge that can't be directly related to the improvement of the economic situation or working abilities of the individual who acquires it, and only with difficulty be compared with the utility derived from economic factors. In other words, it has bad asset quality.

But it does exist, and encompasses knowledge that can sate the person's curiosity, or sustain the person's need for a sense of belonging or appreciation: the means for attending to the social and psychological needs of the individual. Those things that have long been called "priceless".

When the organization knows that it is benefiting from the acquired knowledge but can't identify the mechanism with clarity, and it therefore can't find a reliable way to measure

and value it, we speak of “indirect value”. The organization knows that it is there, but will have serious difficulty gauging the return on investment, and so the utility of any resource dedicated to acquiring this type of knowledge. This doesn’t mean that it can’t or won’t act: advertising is another classic source of this type of utility.

3.4 Axes and intersection: degree, not quantity

With the previous definitions, we can now draw two axes: one that will serve us to track organizational attitude to knowledge, and one that will address the attitude of individuals.




Figure 2

Both axes measure the degree in which the specified knowledge is objective/direct or subjective/indirect. There is no accumulation when we judge a quality, but a degree that is different for every single individual and organization. Just as markets work by aggregation, we can only work by aggregating the perspectives of groups of users: the wider the aggregation, the less meaningful. This will prove essential when attempting to capture value.

The intersection of both axes will therefore need to be in the middle, where value is of a doubtful nature: the point where it ceases to be objective/direct and becomes subjective/indirect. Another way to define that point is where the knowledge evaluated can no longer be considered a proper asset according to the characteristics listed above.

3.5 The Quadrant

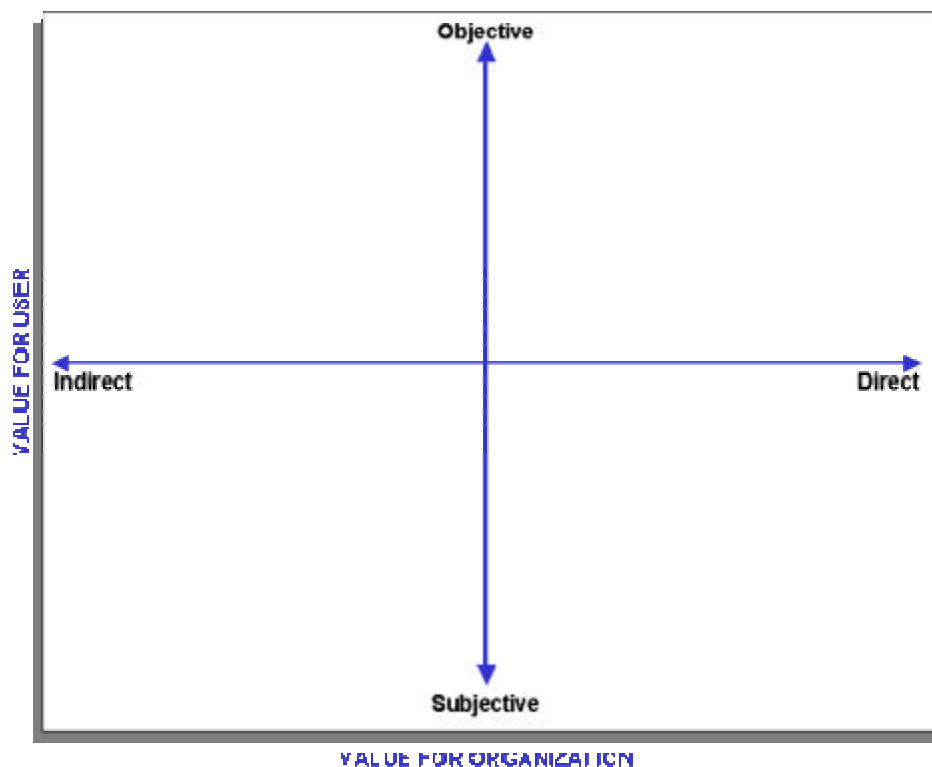


Figure 3

5. Types of Community on the Quadrant

Within the “Knowledge Community” concept there are many different incarnations, with substantial differences among them. Some of these types of Community are quite well-known and studied already.

We will now proceed to place them on the map offered by the Macuarium Value Quadrant, attending to the type of value that is most frequently accrued by their members and the organizations supporting them. Necessarily, this representation will be an area: it could only be more precise by limiting the number of users and organizations.

5.1 Communities of Practice

Communities of Practice (CoPs) are Communities of Knowledge where participants share a common professional interest. On the model, we can point them out as those in which individuals (users) are enabled in a productive way. Thus they derive utility from the potential or actual improvement of their productive capabilities and economic situation, generated through the interactions of the CoP. These communities present the user with an “objective value”.

This Quadrant allows us to underscore something else important: the “practitioners” are individuals, the Community of Practice need not belong to any particular organization for it to accrue benefits from the improvement of its members.

When the organization is able to pinpoint the direct relationship between that knowledge and improvements in its processes and results, we will be talking about “direct value”. When it is unable to discern the way and scale in which this improvement takes place, but is certain that somehow it profits from the CoP, we will be talking about “indirect value”.

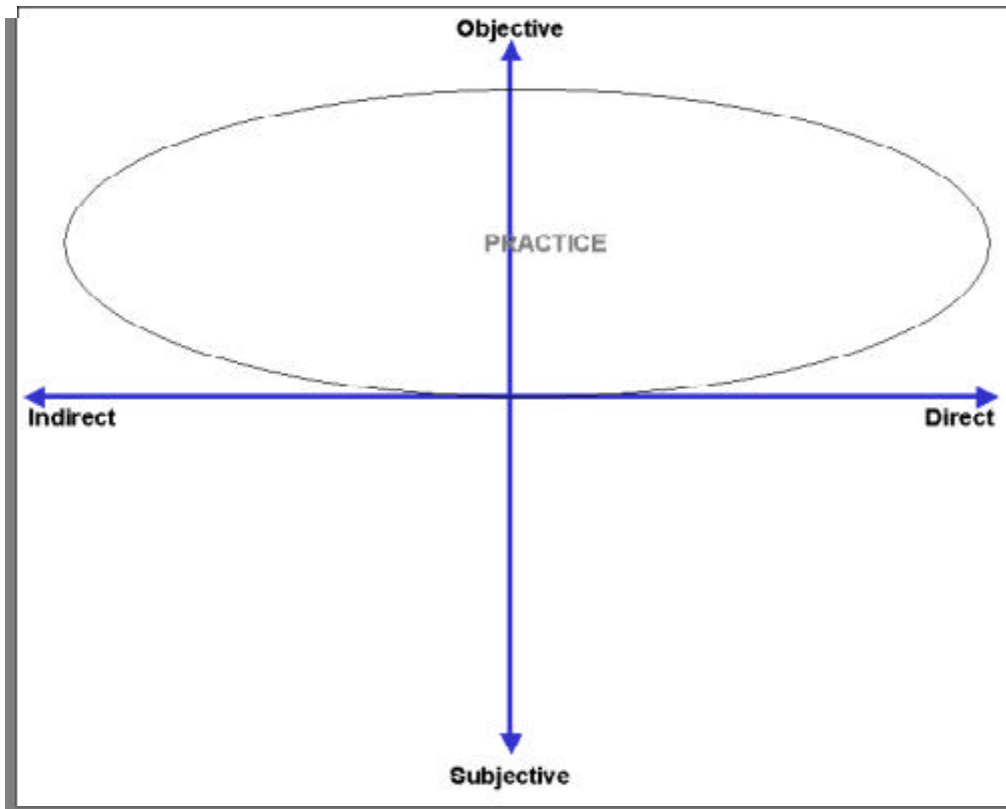


Figure 4

5.2 Communities of Interest

In contrast to the above mentioned type, the variant of Knowledge Communities is defined by a different sort of utility: that derived from satisfying needs such as curiosity, debate, belonging, recognition. In the model, we see that described as “subjective value” for the individual.

For the organization, we find the same division that we saw with the CoPs: a Community of interest can generate either “direct” or “indirect” value depending on how measurable that utility is. That will usually depend in the organization’s relationship with the users’ interest: if it is a supplier, for instance, or a victim, or can exploit the results of the users’ activity, the value can be most direct.

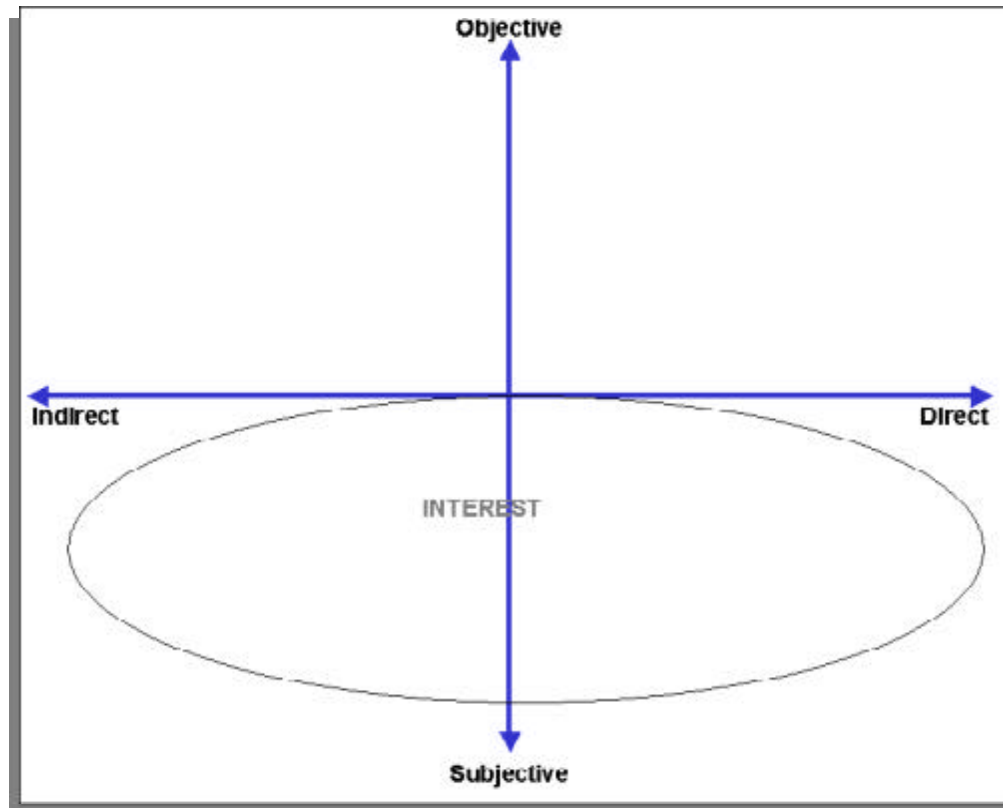


Figure 5

5.3 Amorphous Communities

Reality holds many examples of environments in which knowledge is somehow exchanged, but where the supposed virtues of Communities are hard to find. We see fragmented environments, discontinuous participation, lack of a stable funding mechanism... and we even doubt that these things can be called a Community.

But they can. They can even be fitted onto the Quadrant.

The reason for these apparently failed communities lies squarely in the motivations (goals) that are supposed to push it forward, or lack thereof.

When the individual users derive from them a purely “subjective value”, with no translation into objective value (no cost and no benefit), and the organizations involved receive a totally “indirect value” such as that usually derived from marketing and promotional tools, we have those sorry but frequent examples of “Amorphous Communities”.

Users won't be held to rules (management or structure) nor go out of their way to improve or maintain them. The organizations can't justify serious investment in moderation or IT infrastructure. The community just drifts on, receiving disconnected bursts of activity, and frequently infected by other (cuckoo or parasitic) communities.

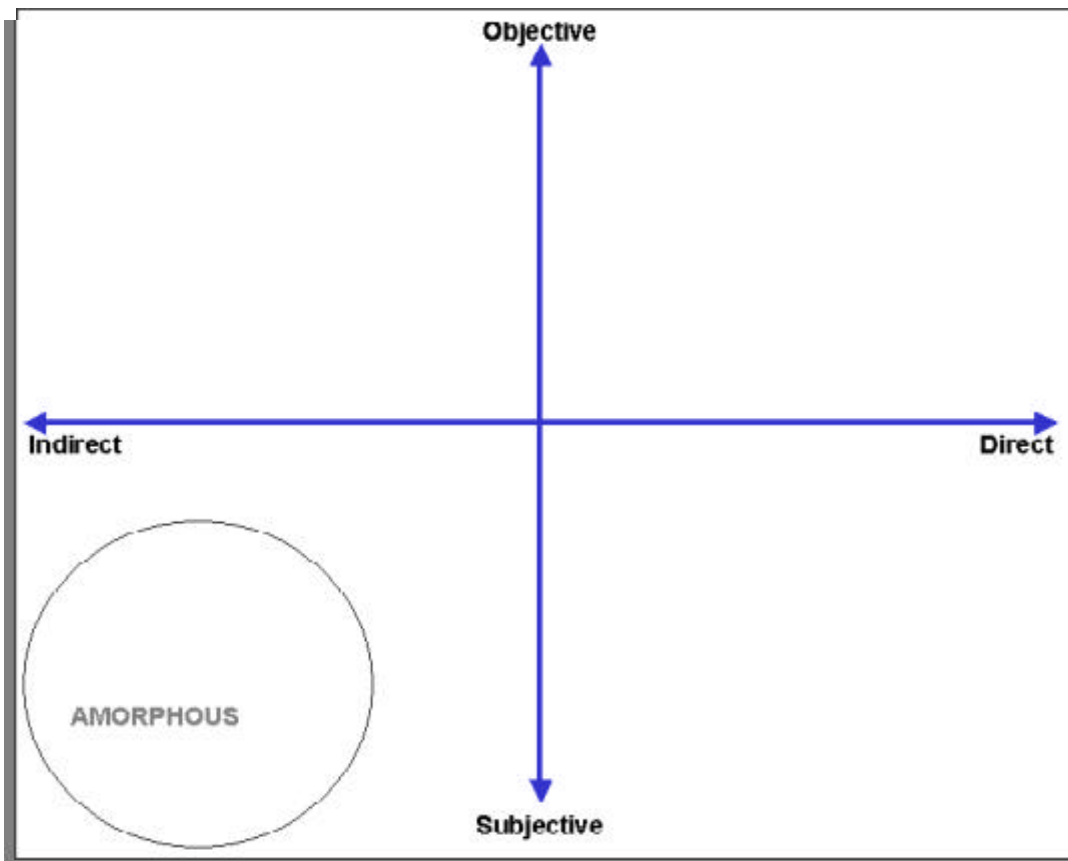


Figure 6

5.4 Project Communities

There is a particular type of knowledge exchange environment that was traditionally hard to place: the one defined by project-oriented teams and knowledge units. Typical examples include product development teams, project implementation teams, customer pre-sales teams...

These groups were often excluded from the definition of Community due to their unusual working methods and organization. We will see that these are nothing more than the results (management) of an extreme set of motivations (goals)... which are present in many other communities.

This can also help us understand the place of Communities inside the organization.

Project Communities are those in which individuals derive an extremely “objective value”, directly related to their jobs and their economic subsistence. It is part of their job, a frequently essential tool. Also, it can be the main way of showing their excellence at work and other capabilities. The behavior of the users toward the Community is directly marked by this perception.

On the other hand, organizations extract a clearly “direct value” from their work. They know every input, and they can control and cost them. They can also relate it with the efficacy of the organization’s processes and the achievement of its aims.

To put it in a nutshell, both parts of the deal regard this type of communities with the utmost seriousness and want the highest control on them. This has a direct reflection on the way they are managed and organized.




Figure 7

6. The value and the capture: open discussion

So we arrive at the final stage: now that we have established that knowledge may generate different types of value through the participation of people in different types of Community (which is anyway intuitive) we come to the hurdle that has kept this phenomenon from full recognition by the business world. Can Communities capture the value they generate? Can they be a worthwhile investment for their promoters?

By “capturing”, as we explained above, we indicate that the Community manages to keep a proportion of the value it created for third parties: in other words, that it somehow manages to get paid for it.

Due to the nature of Communities (they derive value from the participation of members) this has proven difficult to manage: charging for use went against the perceived goal of increasing the user base. And it was usually found difficult to charge organizations due to the lack of a way to prove a value-creating quality worth funding. We will now try to see how this can be done.

As a general rule-of-thumb, let’s say that charging for something usually depends on the ability to withhold it... or at least, the most valuable portion of it.

6.1 Who pays for Communities of Interest

From the standpoint of individual members, a Community of Interest is a source of information and (frequently) of social integration. Therefore, the promoters should strive to find ways of maximizing these, and charge for their use.

Typical examples will be: access to better and deeper search tools, to member galleries and directories, to private “inner circle” forums and services, even small tokens as the use of an identifying picture or “avatar” on the forums. The rate of uptake will not be great, but it will depend more on the need being satisfied than on the tool used. Pricing will most frequently be low, since value is subjective, but production costs for this type of services need not be high, and it is feasible to extract a profit.

From the standpoint of organizations, a Community of Interest is a natural gathering of people with the same interests and (usually) similar demographics. A well-selected Community of Interest is a prime marketing target, not just for advertisement but also for customer study.

It is frequent to see undefended Communities bombarded with spam and publicity: see the Usenet news or any low-administered forum. But a serious Community built on private premises can ward these advertisers off, and channel them in more useful ways: paid advertising, such as banners, text-ads, infomercials, product reviews, search-related links, and more advanced services. All of which, of course, can be charged for.

Any other ideas, examples and cases will be very interesting to hear.

6.2 Who pays for Communities of Practice

Communities of Practice may be internal (belong, in one way or another, to the same organization than its members) or external (members participate without regard for their respective organizations).

From the standpoint of individual members, a Community of Practice is a gathering of fellow practitioners in a professional subject. It is a place at which to solve practical doubts and gain recognition as an expert (frequently the best way for career improvement). As before, the goal should be to find ways of maximizing these, and charge for their use.

When the Community is already a recognized one, and there are few or no free alternatives, it is not absurd to charge for membership. Equally, it is feasible to charge for access to older messages or advanced search features, since they can very often provide really valuable insights into problems. More advanced ways to capture value include a listing of experts (to be included in), and even a hotline (to address those –or other- experts at a given rate, any time of day). The uptake should depend on the real ability of the CoP to solve actual production problems. And the pricing can be carefully graded, up to really high-value services.

From the standpoint of organizations, an external Community of Practice is a natural gathering of customers (or, sometimes, providers). In addition to being a prime marketing target, with the added interest of being critical –bad opinion garnered there can kill sales easily, and viceversa- it also works as the best and cheapest way to provide support for users of the organization’s products and services: the members support each other, and since the real experts usually are the product’s most frequent users, they will often give better support than an organization’s hotlines.

Those marketing interests can be channelled as explained for Communities of Interest; the support element lends itself to a different possibility: the creation of sponsored communities, even monographic or dedicated ones, in which the organization’s product has center stage and brooks no destructive criticism. Simply offering external support can be a tempting (and valuable) service; those more dedicated options are high-value possibilities, although they are necessarily limited in number and hard to manage: they may comprise a Community’s source of value as an objective judge.

From the standpoint of organizations, an internal Community of Practice can be a lot of things. CoPs can be used as enablement tools to enhance knowledge transmission and work

efficiency between a group of knowledge workers. It may be used as a steward for a range of knowledge domains, or even upkeeper of knowledge objects. It may be an emergency-management tool. It may be a talking-shop and decision-making centre. But this all depends on the motivation of the CoP, and the role the organization is prepared to let it play.

In essence, the ways to make the company pay are these same that figure as the goals for the CoP: if those are valuable for the organization and it has no alternative but to fund the CoP or lose them, it will pay. The problem will usually be in creating the decision point and making clear the value of the goal. The “price” thus captured will often only consist of tacit member permission to participate and some technical means to articulate the CoPs.

Any other ideas, examples and cases will be very interesting to hear.

6.3 Who pays for Amorphous Communities

From the standpoint of individual members, an Amorphous Community is a source of gossip and loose social integration. Therefore, the promoters should strive to find ways of maximizing these, and charge for their use.

Typical examples will be: access to private forums or chat rooms for groups of friends, gaming clans or other associations within the amorphous community, and related services such as identification tools (even small tokens as the use of an identifying picture or “avatar”, signatures...), and self-expression tools such as personal pages or “blogs” (weblogs). The rate of uptake will not be great, since the market is crowded and there is a lot of offer. Pricing will most frequently be low, since value is subjective, but production costs for this type of services need not be high, and it is feasible to extract a profit.

From the standpoint of organizations, an Amorphous Community is a stochastic gathering: there is no defining reason for all of the members to be there. It is only fit for wide-focus (and low-paid) marketing initiatives. For it to become an interesting marketing target, it needs to be parcelled into groups of like-minded (or like-buying) individuals. The Community administrators can frequently do this through the registration and mining of user profiles. The traffic of these profiles (including even email data) can be a source of low-level income too, if needed.

Any other ideas, examples and cases will be very interesting to hear.

6.4 Who pays for Project Communities

Project Communities may be seen as an extreme example of Communities of Practice, in which members are also usually integrated in a formal relationship for the achievement of a shared professional goal. They are usually internal to an organization (or to a joint-venture, or other closed collaboration initiative), but may eventually use external resources: for instance, a group of e-lancers working together on a project using a remote collaboration tool, or a group of Open Source developers refining code at a specialist site like SourceForge.

Project Communities also have a habit of including even more face-to-face interaction than the rest of community types: it is not unusual to see examples in which it is the most frequent form of relationship, and the electronic support means serve only to exchange materials, references, schedules and messaging.

From the standpoint of individual members, a Project Community is a working tool. It provides ease of coordination, a better flow of work, and increased efficiency. Unless they work for themselves or are in any way managers of their own resources, they will not be able to buy or hire the tool.

On the other hand, e-lancers or independent professionals may see fit to hire the technical and professional means to set up a Project Community for the duration of a project.

From the standpoint of organizations, a Project Community is essentially a tool for its employees: it enables projects to develop in a more efficient way. As such, they must be able to measurably prove that fact through reports and analysis. The value accrued and thus proven will be the key to calculate the investment that the organization will rationally be prepared to provide for the maintenance and development of the Project Communities infrastructure.

Any other ideas, examples and cases will be very interesting to hear.

7. Notes and references

- [1] We will go into the measurement and expression of the value derived from corporate intellectual capital in an upcoming paper. Check at <http://www.macuarium.com/econoclasta/network> for it after July 2003
- [2] "Views of Knowledge are Human Views", by G. Duek, which can be read on the web at <http://www.research.ibm.com/journal/sj/404/dueck.html>
- [3] We are assuming the classical economic models of "homo economicus", and the utility or indifference curves: further information may be found http://www.wikipedia.org/wiki/Homo_economicus and http://www.wikipedia.org/wiki/Indifference_curve
- [4] On the mechanics of price formation, see http://www.wikipedia.org/wiki/Supply_and_demand
- [5] By "private", we simply mean that a particular entity is the owner of those rights. The legal quality or type of owner is irrelevant here