

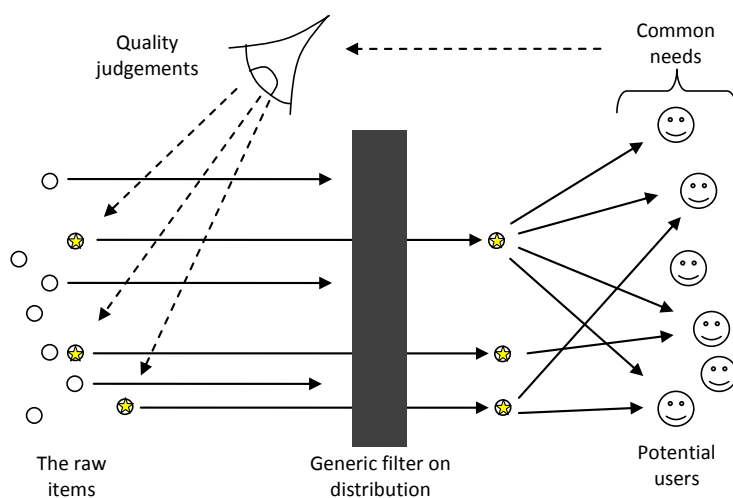
## Disaggregating Quality Judgements

*A Position paper for the Quality Commons meeting, Paris 2010*

*Bruce Edmonds, Centre for Policy Modelling, Manchester Metropolitan University*

Something has *quality* if it might be what you want, but this can never be the whole story unless your needs are identical with those of others. Quality is a rough proxy for an item's suitability with respect to your needs, based on the fact that your needs will have a lot in common with those of others. If everybody's needs are different, for example needing the right word to express something, then quality has no meaning – there is no such thing as a good (or bad) quality word. If there is a substantial commonality between the needs of different people, for example that an apple is ripe and not gone bad, then an object that fulfils what is common to those people's needs can be judged as to its quality – thus one can have bad apples. I will call selecting objects as to their quality the *quality heuristic*.

The quality heuristic goes like this. A person, persons or process make judgements on a set of items as to their quality. All those searching among these items for something that meets their needs can reduce their search cost by concentrating on those judged as having higher quality. A version of this is illustrated in Figure 1.



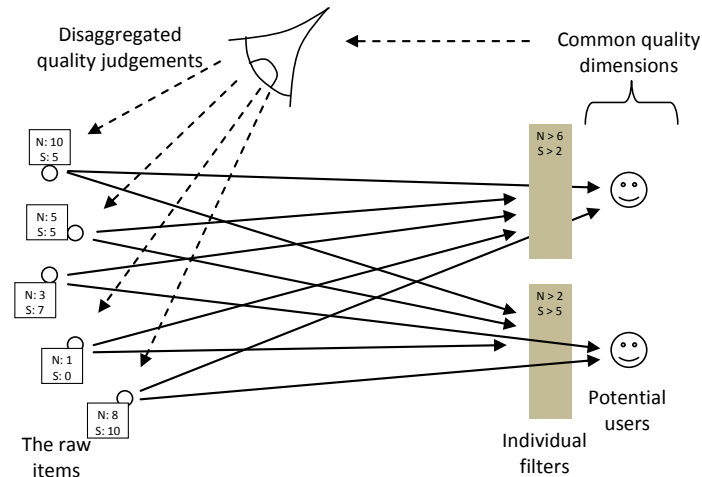
**Figure 1. The Quality Heuristic (using a binary filter on distribution)**

The quality heuristic becomes important when it is costly or impossible to find an object that meets your needs yourself. For example, if only a few items can be made available for many to consider then it is important that those items that are selected meet the needs of as many of the potential users as possible. This was the case when publishing an academic paper and distributing it was costly – only a few papers could be published so it was important that only quality papers were. Such quality judgements were possible because there was a lot in common between the academic needs of the readers – they all wanted similar things, such as: soundness, importance, novelty, verifiability etc. Thus bottlenecks to finding and accessing what one needs makes the quality heuristic necessary.

Of course, there may be only a partial alignment between the needs of potential users. So in the apple example, although almost nobody wants an apple that is not ripe or has gone bad, some will want an apple that goes well with their cheese, some a small sweet apple for a child's lunchbox, and some an apple for making an apple sauce. The more closely that a quality judgement is restricted in context so that the needs of users in that context coincide, then the more useful (and hence more meaningful) a quality judgement can be. Thus it makes more sense to talk about an excellent cooking apple than simply an excellent apple (unless the context of use has been implicitly or explicitly specified). The greater the commonality of needs, the more effective the quality heuristic is. For a system that produces a binary or single dimension of quality judgement then the commonality between the users is critical to its success, it works depending on whether what is suggested by one member of the community is actually what another member wants. However when one disaggregates the quality judgements this becomes much less of a problem.

With respect to academic information, there is no longer a great cost to the distribution of papers (Harnad 1998). However, users still have a difficulty in finding the information one wants due to the sheer quantity of papers available – it is the search cost that is the bottleneck now. Thus there is still a need for quality information, but due to the different bottleneck, there are better solutions than the one illustrated in Figure 1. The users still need access to the

information from quality judgements, but there is no need to impose a filter on the distribution of information. Rather if rich quality judgement on the items can be associated with them and made available to users (say in a database of judgements plus link to the item) then the users can use this information in a flexible way to find those items that meet their needs. In this way the needs of users can be met in a more accurate manner, dissolving binary quality judgements into a richer and more informative set of judgements with respect to different aspects of the items. Thus if judgments as to a paper's novelty, importance, completeness, clarity etc. were captured, then one user could search for all items that were very novel but with only a minimum level of clarity whilst another could insist on high completeness but not need any novelty. Monolithic quality judgements that averaged *all* the potential readers needs would be replaced by a more flexible system, but one where quality judgements are still central.



**Figure 2. Delaying selection using rich quality information**

In this way items are tagged with extra, judgmental information as to their qualities, that allow the users to solve the search problem, but in a way more specific to their needs. The binary publish/don't decision is replaced by that of making quality judgements along many dimensions. This need not be done by one person or institution, and could be decentralised. Indeed the user could select which *kinds* of sources of judgemental information they will use. Thus social simulators might choose to only utilise judgements by a specially formed panel of reviewers, from any other social simulators, or indeed from anyone at all. There is nothing to stop the providers of judgemental information being judged in a similar way to the items that are judged. One need only have to trust a source of judgements as to the trustworthiness of reviewers and that could allow one to limit the sources of judgemental information used.

To make such a system effective: there needs to be a protocol for encoding the quality judgements in the various dimensions; a sufficient amount of quality judgements need to be captured and stored in an accessible manner; the search engines need to be able to utilise this information so as to allow for quality-enhanced searches. The process of bootstrapping trust in the system is essential to its uptake by users, the amount and soundness of the quality information necessary to make the system at all useful. A feature built into a browser could allow a pre-registered reviewer to simply set sliders as appropriate to the paper they are viewing in the browser, optionally enter a comment, and press submit. One could combine the output of Google with the judgemental information to access the results, optionally browsing attached comments of the results, before deciding to whether to read the paper. Google has introduced a facility that allows the public annotation of any web page, called "Google SideWiki" (<http://www.google.com/sidewiki>). Given an agreed structure this could be in a form that could be (1) written as above using review feature and (2) combined with search results, simply removing those items that positively failed on the given quality criteria. I first suggested this idea in (Edmonds 2000).

It should be noted that the efficient selection and provision of papers to readers is only part of the purpose of a traditional, distribution filter it also has social aspects. A filter on papers may help the coherence and focus of a field, by ensuring all its members read a similar body of work, it may be a filter against dissent and criticism of a field, and it may act in political ways allowing a collective determination of future directions and issues.

Edmonds, B. (2000). A Proposal for the Establishment of Review Boards - a flexible approach to the selection of academic knowledge. *Journal of Electronic Publishing*, 5(4). <http://www.press.umich.edu/jep/05-04/edmonds.html>

Harnad, S. (1998) On-Line Journals and Financial Fire-Walls. *Nature* 395(6698): 127-128. <http://www.cogsci.soton.ac.uk/~harnad/nature.html>