

Going Meta: Norm Formation and Enforcement in the Stack Exchange Network

Paul Matthews
UWE Bristol
Coldharbour Lane, Bristol
+44 117 32 83353
paul2.matthews@uwe.ac.uk

This paper presents research into the way norms are developed, expressed and enforced on sites that are part of the Stack Exchange (SE) social question-answering network. This network has a number of topical knowledge exchange communities using similar underlying software, enabling a focus on variation in social design. SE also separates community-related discussion from topic-specific content through the use of its “Meta” sub-sites. These were analysed together with their main sites for variation in the development and enforcement of norms. Norms expressed through explicit community policies seem rather less important than those embodied in busy discussion threads on the Meta sites. While Meta participation was fairly uniform across communities, different emphasis on scope and quality led to variation in Meta discussion and the way that norms were enacted through question closures. The social distribution of moderation work was also uneven between sites, with some sites having a few highly active moderators involved in question closure. The level of closures across the sites studied did not seem to significantly discourage participation. Indeed, modelling the effect of closures on quality and engagement indicated that low levels of closure enable “legitimate peripheral participation”, the process by which newcomers can become inducted and make contributions of increasing quality over time.

Online communities, norms, moderation, social question-answering, SQA

"The ideal moderator does as little as possible. But those little actions may be powerful and highly concentrated. Judiciously limiting your use of moderator powers to selectively prune and guide the community — now that's the true art of moderation." - Jeff Atwood

This ongoing research is investigating the social design and the development and implementation of norms on the family of question-answering sites belonging to the Stack Exchange (SE) network. These SE knowledge exchange sites are of particular interest because 1) they all use the same or very similar underlying software and so the technical elements are somewhat controlled, and 2) they usefully segregate the domain-specific question and answer content from the “meta” discussion about how to organise and run the community.

All of this latter discussion is placed on a sub-site of the community to which members have access – the “Meta” site.

These two features make the SE community a useful candidate of study in order to better understand how social design can vary between communities of interest. Particularly, we can look at the methods by which the different communities develop norms of behavior and how these norms are subsequently enacted.

The initial work presented here focusses on question closure, as this can be considered a strong community sanction in enforcing the quality standards they have collaboratively developed. On SE sites, closures occur when five elected moderators of high standing in the community agree that a question does not meet community guidelines.

Social norms have been theorized in several ways. We can distinguish injunctive and descriptive norms, to describe what others approve of and what they actually do respectively [3]. The injunctive norm may vary from the descriptive and hence reduce the overall power of the norms' effects. A similar bimodal concept is that of explicit or implicit norms in online communities with explicit norms being, for instance, reified in FAQs and implicit norms more connected with the overall “tone” of debate [1].

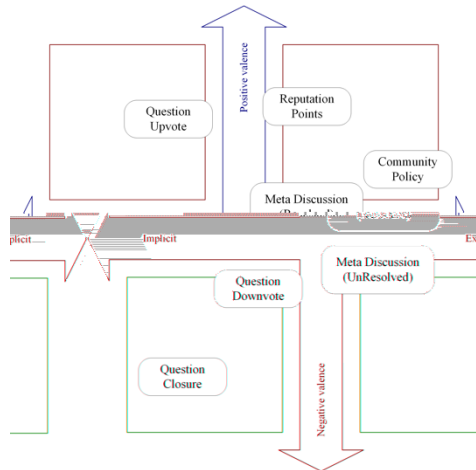
In Burnett's theory of normative behaviour, norms are enacted through four mechanisms: (1) social norms; (2) worldview; (3) social types; and (4) information behaviour, which might be summarised respectively as: a shared sense of rightness and wrongness; a collective sense of what is important; a set of signature behaviours (more or less constructive) and the extent to which individuals seek out and take on board information. Burnett offers case studies in online communities in which each of these can be distinguished [1].

Norm propagation may be vertical, oblique or horizontal [4]. The first two are associated with the transmission of norms from authority figures to those of lower status, the last with transmission between peers – both of these are at play in communities such as

how knowledge of norms and the ability to cite them led to more successful arguments with respect to the deletion of pages [5] “*This shows Wikipedia’s focus on precedent and rules, and the tendency to discuss articles in terms of both values or community norms and article contributors or supporters*”. The authors note a distinct difference in perception, with newcomers more likely to take page deletion personally, but experienced editors showing more rational argument [5]

In SE communities, explicit and implicit norms may be enacted in a number of ways (Figure 1). The most explicit norms are articulated in the site help pages. Thereafter, the discussion on the Meta sites is related to the interpretation of these policies and the identification of the need for new guidance. This may be more or less explicit, given that some debates are resolved through participant consensus, whereas others are more polarized. On the more implicit side, users can “horizontally” communicate norms through voting questions up or down and through commenting on them. At the most extreme, moderators can vote to close questions considered unsuitable. This is probably one of the most extreme forms of normative behaviour visible on the sites.

Both implicit and explicit norm enactments have a valence in terms of the emotional impact on the subject or originator. There is an assumption that for a healthy community, and to encourage the participation of newcomers, negative valence actions need to be limited or mitigated [2]. This may be related also to theories connected to social dynamics within communities of practice, where a diverse mixture of beginners and “old-timers” enables long-term sustainability [6].



This research is using the StackExchange API¹ to extract activity in the form of basic site metadata and activity, questions / question closures and user comments from a selection of StackExchange topic sites listed in

Table 1.

Sites covering a range of contrasting topics were deliberately chosen to highlight contrast and increase the likelihood of distinct user membership.

¹ <https://api.stackexchange.com/>


To analyse close behaviour, samples of 30 questions from 10 random dates were extracted for 2012, 2013 and 2014.

| | | | | | |
|---------------------------|---|------|--------|-----|--------|
| MathOverflow | professional mathematicians | 2009 | 67,000 | 4 | 51,000 |
| English Language & Usage | linguists, etymologists, and English language specialists | 2010 | 22,000 | 8 | 21,000 |
| Science Fiction & Fantasy | science fiction and fantasy enthusiasts | 2011 | 28,000 | 8.5 | 41,000 |
| Android Enthusiasts | enthusiasts and power users of the Android OS | 2011 | 36,000 | 8.5 | 94,000 |
| Philosophy | those interested in logical reasoning | 2011 | 5,600 | 16 | 13,000 |
| Seasoned Advice (Cooking) | professional and amateur chefs | 2010 | 14,000 | 6 | 27,000 |

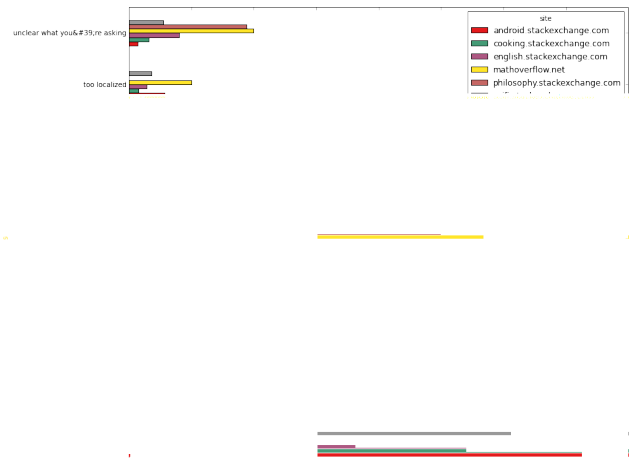
To further investigate the potential effect of the closure sanction on community dynamics, a model was built of a community asking questions over time. The model set-up and parameters are further described below.

The proportion of users participating in both the main and Meta sites varies from 5 to 15%, providing a discernable linear relationship between the user base of the main community and the Meta section (Figure 2).

Actual activity, in terms of question and answer volume is rather lower with question volume varying from 1 to 7% of the main sites. MathOverflow and Android were least active, Sci-fi and Philosophy the most.



This again revealed differences in the application of norms, with more questions closed for being off-topic in MathOverflow than other sites. Interestingly, the Philosophy site applied the “not constructive” close reason more frequently – perhaps understandable given the nature of philosophical debate.




One of the canonical sources for scope is the “What topics can I ask about here?” help page, which unlike others, varies between StackExchange sites and can be edited by the site moderators. These pages were downloaded and compared. The longest page, at 3122 characters (removing navigational and static page elements) belonged to English. In contrast, the shortest page, at half this size, belonged to cooking (1595 chars), followed by science fiction (1645).

A Levenshtein pairwise comparison of the pages’ content revealed Android and Cooking to be most similar (0.60), with Cooking and English the least (0.49). This would seem to be largely an artifact of the different lengths of the pages. That said, the results together indicate the English help page has had the most editing effort applied. The more similar pages have clearly been simply appropriated from “boiler plate” text.

To identify substantive discussion on the Meta sites relating to norms, questions were selected by tag over the period 2012-2014. This data indicated large variation in application of the tags, though it is not yet clear if this indicates a difference in community makeup. Certainly, the focus on etiquette in Math Overflow is of interest in contrast to the attention given to close reasons and allowed topics on the Sci-Fi community. We might hypothesise more consensus among the MathOverflow community as to the scope of topics considered valid (and it is certainly an older community).


Do closures discourage participation? This was investigated by comparing the pooled data on a user basis - those users experiencing closures with those never experiencing them (users who were active in the sample period). From the resulting sample of 200 users in each group, the answer seemed to be that users were only marginally more likely to post more if they didn’t experience closures (Figure 4). Indeed in the non-closure group, a great many users still only posted a single question. This remains a coarse analysis and it may be instructive to see if there is any variation between sites in future work.

| | | | | | |
|--|----|----|----|-----|----|
| | | | | | |
| | 20 | 7 | 1 | 5 | 5 |
| | 11 | 4 | 1 | 14 | 5 |
| | 10 | 2 | 0 | 22 | 11 |
| | 43 | 37 | 3 | 122 | 20 |
| | 83 | 11 | 27 | 38 | 14 |
| | 77 | 29 | 5 | 0 | 4 |



To investigate the treatment of closures between sites, closure reasons given in the three-year sample were compiled and normalized to a percentage of all closures. Results are shown in Figure 3.

To investigate the pattern of moderation around closures, the users associated with closure votes in the sample period were collated (Figure 5). These indicated a greater diversity in moderator activity on some communities (Sci-fi and Math) compared with others, where there were more prolific individuals (Cooking, Android). While all community moderators are contributing in their own time, this does indicate a between site variation in the distribution of norm enforcement among moderators.



Based on realistic SE community attributes for engagement and question closure rates we developed an agent-based model to compare the impact of question closure on quality and collective expertise. By “expertise” here we mean a combination of subject-matter expertise and question-asking expertise, the latter reflecting a high level of induction into community norms.

The following assumptions were made concerning community makeup and dynamics:

1. The community starts with 100 members. Members start with 1 reputation point and are randomly assigned a “questioncraft” attribute ranging from 0 to 0.5
2. At each time interval, members ask a question with probability:

$$p_{ask} = \left(a + \frac{b(n_{asked} + 1)}{t_{member}} \right) \left(1 - \frac{r_{closed}}{c} \right)$$


Where r is the closure rate (closures as a fraction of total questions), n is the number of questions previously asked, and t is the length of membership in time periods. a , b and c are constants set to 0.8, 0.4 and 3 respectively.

The intention was to approximate the three intuitions that: (i) members who have asked before are more likely to ask again; (ii) new members are more likely to ask a question; and (iii) the closure rate has a small inhibiting effect on the likelihood of a member asking a question, for fear of sanction.

3. If the probability of asking exceeds a threshold (0.8 in our model), then users ask a question that has a quality relating to their questioncraft and experience.
4. We set a closure threshold for question quality. If the question asked falls below this, the question is closed and the user receives a negative reputation adjustment weighted to their question quality. If it is of acceptable quality, then the user receives positive reputation points, again proportional to their question’s quality.
5. All surviving users who asked a question (users with a negative reputation score leave the community) receive a small increment to their questioncraft.
6. At the end of the time period, new users join the community. Here, we add 10 users with an adjustment according to the closure rate. The higher the rate, we assume, the less “friendly” the community appears to new members.

A simulation using the same initial questioncraft distribution was run for 100 time intervals for four closure thresholds (0.02, 0.05, 0.10 and 0.15), reflecting the range of closure levels seen in the stack overflow communities (see Table 1).

Figure 6 shows results from these four runs. We find similar mean question quality at each closure level, though the range is more variable at lower closure rates. With lower closure rates, we also find a slightly smoother reputation distribution.



The SE network has given rise to successful, knowledge-rich communities, due largely to a focus on community self-organisation and moderation. The preliminary work presented here indicates a broadly similar level of participation in the Meta sites of different interest groups, where the formulation and discussion of norms occurs that are then enacted in the main site.

Some communities seem to have more of a problem with scoping – deciding what exactly is on-topic for their community. Conversely, where a specific and agreed scope obtains, sanctions in the form of question closures are more notable for scope-related reasons.

Closures on initial inspection are selectively applied and balanced and do not seem to adversely affect participation. This is interesting, as all SE sites are struggling with striking the correct balance between inclusivity and quality. Our modelling suggests that high closure levels tend to reduce variance in quality, at the expense of very high quality as well as low. This may explain why the closure rate remains low in some of the more successful SE communities.

While closures have been taken as an implicit norm, they also act explicitly, because close questions remain visible after closure with the reason given. This may have a larger effect on community behaviour than on the original poster themselves.

The findings from closure modelling appear to tally with the dynamics of Wenger's "legitimate peripheral participation" [6]. A low closure threshold encourages newcomers to ask questions and through participation they are inducted into the mechanisms of community quality. With more censorship, it becomes less likely that newcomers have a chance to accumulate experience and reputation and a smaller set of longer-term members come to dominate the high reputation cohort.

This work has exposed some of the more explicit and readily measurable norm-related actions. A clear target for further investigation are the discussions around interpretation and application of norms for different communities, where we have already seen there are differences in focus (and controversy). These will probably need to be investigated part-qualitatively, though initial work using semantic analysis is being trialed.

The study deliberately compared contrasting site topics. It would also be revealing to compare sites with more closely related subjects (e.g. several SE sites are technology-oriented). A hypothesis here might be that the differences would be less marked.

Variations in norms between sites need to be linked to outcome measures, including Sense of Virtual Community and resource quality. This will give a better indication of where particular norm calibrations lead to successful social designs.

[1] Burnett, G., Besant, M. and Chatman, E. A. 2001. Small worlds: Normative behavior in virtual communities and feminist bookselling. *J. Am. Soc. Inf. Sci. Technol.*, 52, 7, 536-547. DOI=10.1002/asi.1102.

[2] Petrič, G. and Petrovčič, A. 2014. Elements of the management of norms and their effects on the sense of virtual community. *Online Information Review*, 38, 3 (04/29; 2015/02 2014), 436-454. DOI=10.1108/OIR-04-2013-0083.

[3] Cialdini, R., Reno, R. and Kallgren, A. 1990. A Focus Theory of Normative Conduct: Recycling the Concept of Norms to Reduce Littering in Public Places. *Journal of Personality and Social Psychology* 58, 6: 1015–1026.

[4] Savarimuthu, B. T. R., Purvis, M., Purvis, M. and Cranefield, S. 2009. Social Norm Emergence in Virtual Agent Societies. *Declarative Agent Languages and Technologies IV*, 5397, 18-28.

[5] Schneider, J., Samp, K., Passant, A. and Decker, S. 2013. Arguments About Deletion: How Experience Improves the Acceptability of Arguments in Ad-hoc Online Task Groups. In *Anonymous Proceedings of the 2013 Conference on Computer Supported Cooperative Work*. (San Antonio, Texas, USA,). ACM, New York, NY, USA, 2013, 1069-1080.

[6] Wenger, E. (1998) *Communities of Practice: Learning, Meaning, and Identity*. Cambridge University Press, Cambridge, UK.