Paperless Philosophy as a Philosophical Method

David Bourget

June 3, 2010

Abstract

I discuss the prospects for new forms of professional communication in philosophy. I argue that online discussions and online surveys ought to play a more important role in communications between philosophers than they play today. However, there are major obstacles to the widespread adoption of these media as channels of communication between academics. I offer an overview of these obstacles and sketch a strategy for surmounting them. The strategy I propose involves the development of a new kind of service which could expand the reach of the analytic method in philosophy.

A philosopher's work centers on the production, delivery, and consumption of books and articles. New communication technologies have enabled faster, more convenient means of contributing to and drawing from the philosophical corpus, but there has been no major shift in the format of research communications since the foundation of the Royal Society in the 17th Century. Philosophers still communicate their research mostly through books and articles, whether in print or delivered verbally through talks and lectures.

In this paper, I discuss the prospects for new forms of professional communication in philosophy. I am particularly interested in the future of online discussion spaces and online surveys; I believe that these types of service ought to play a more important role in communications between philosophers than they play today. I will explain what potential I see in these services and offer a tentative diagnosis of why their potential remains largely unrealized. Having identified some problems with current discussion spaces and survey services, I will sketch a new kind of online service which integrates discussion spaces and surveys in a way which could help overcome their individual shortcomings. The service I propose could expand the reach of the analytic method in philosophy.

1 Online discussion spaces: potential and reality

By "online discussion spaces", I mean all online spaces which enable participants to pursue public (or semi-public) conversations with groups of individuals asynchronously. This definition encompasses many different kinds of online service, including discussion forums, weblogs, and the kinds of messaging spaces found on social networking services.

Online discussion spaces have much unrealized potential as a medium for professional communications between philosophers. As anyone who regularly attends philosophy seminars will know, sometimes it does not take many words to raise serious doubts about a claim or a whole paper. Likewise, sometimes a claim can be effectively supported in few words, e.g. by citing a relevant experiment. But in today's paper-centric system, many verbally brief points which are potentially just as valuable as whole papers remain largely uncommunicated to the profession as a whole. This happens for nearly all contributions which are a) too small to constitute full papers and b) for which the contributors have no sufficiently good or relevant articles in which to insert them. Online spaces could provide a venue for these "small" contributions. Instead of being confined to a limited number of peers lucky enough to be at the right place at the right time, all such contributions could find a home in online discussion spaces.

The fluid and iterative nature of online discussions also makes them more efficient means of communication than papers for most purposes. When one presents an argument in a talk or paper, one must have a fairly good idea of how one's audience will react to what one says at each step along the way; otherwise one might speak for a long time to no effect. One has to provide enough background information to enable one's target audience to follow, but not too much. One has to carve out the topic in a way that will prove interesting to sufficiently many potential readers or listeners, and one has to plan an argumentative course that will sustain interest. One has to successfully identify the points one needs to invest more energy defending in order to make one's argumentation as effective as possible. And so on. These tasks become increasingly more difficult the longer one's communication goes on without substantive feedback or the possibility of a course correction. This makes papers very difficult to prepare and prone to spectacular failure—one can end up assuming something most currently deny, arguing for a point a large majority already agrees with, failing to communicate anything worthwhile due to insufficient consensus on terminology, etc. These problems are often not philosophical problems but pure communication problems: they are rooted in our ignorance of others' dispositions towards the claims we utter; they are problems of a sort we would not have if we could simply share the contents of our minds with each other through telepathy. I would not be surprised if it turned out that we spend significantly more time and energy on communication difficulties than philosophical problems per se when preparing articles. By contrast, in conversation—or in an online discussion—the possibility of adapting (or abandoning) one's line of argument based on feedback received along the way largely eliminates most of the aforementioned risks and greatly diminishes the effort required to avoid them.

Perhaps the online discussion spaces most used by professional philosophers to discuss philosophy are weblogs. There are currently a large number of weblogs dedicated to specific areas of philosophy, and a number of these have been hosts to high quality exchanges between professional philosophers. But while there are successful weblogs, the extent of the success of weblogs in general is limited compared to their potential or that of online discussion spaces generally. As far as I can tell, a minority of philosophers contribute philosophy (as opposed to news, social commen-

taries, etc.) to weblogs or other discussion spaces. Usage should continue to grow over the years, but there are some fairly obvious reasons why participation in online discussions of philosophy among professionals is likely to remain limited if no major change occurs in these services, the information environment which surrounds them, or our professional culture.

To a first approximation, the main impediment to a large-scale use of online discussion spaces to communicate research outputs is the former's lack of prestige. To make a career as an academic, one must publish in prestigious venues. And the more one publishes in prestigious venues, the more rewarded one is. Since contributions to online discussion spaces carry little prestige, in most cases, most academics are better off reserving their precious research time to work on articles for publication.

A moment's reflection on the nature of prestige shows how deep this problem runs. The prestige of a journal is grounded in the community's perception of the relative worth of its content: a journal is prestigious for a given community to the extent that individuals in the community tend to expect articles published in it to be better than articles not published in it (ceteris paribus). The relevant measure of quality is a complex one which encompasses such factors as relevance, strength of argumentation, clarity, and probably many others—all factors which typically enter into one's judgment of an article as "good" or "bad". Parallel analyses apply to the prestige associated with other entities. For example, a book publisher is prestigious for a given community to the extent that individuals in this community tend to expect books it publishes to be better than books it does not publish (ceteris paribus). A university is prestigious for a given academic community to the extent that individuals in this community tend to expect individuals who are members of the university (students, faculty) to be better academics than individuals who are not (ceteris paribus). Prestige is ultimately grounded in certain kinds of limited resources, e.g. publishing space and academic positions. A resource (and, by extension, the entity that controls it) is prestigious for a community to the extent that someone or something who has acquired it tends to be regarded as better (by some norm relevant to the community) than someone or something who has not, everything else being equal.[1]

As it is, contributions to even the best online discussion spaces carry little prestige for the community as a whole because the community as a whole does not expect them to be of high quality (which is not to say that relatively small communities which use some of these services do not correctly believe them to host high quality discussions). Some discussion spaces carry no or little prestige because contributions are not filtered in any way, which has the consequence that contributions to these venues do not come with any guarantee of quality (and, in practice, often fail to live up to academic standards). Others are filtered, but none are heavily filtered. Filtering also has to be done in a way that is known and trusted by the community as a whole in order to confer prestige, and this is unlikely to ever happen for online discussion spaces moderated by one or a handful of academics in a specialized area following their own rules (the current model for the most successful weblogs). What

---

[1]Note that this account of the nature of prestige is compatible with one's finding that a given prestigious journal generally publishes articles of low quality–what confers prestige is the view of the community as a whole.

is more, it would not be enough for philosophers to find certain online discussion spaces prestigious in order to remove the difficulty: we would also need to convey the prestige of our online discussion spaces to the Dean's office and funding bodies. Otherwise, academics seeking advancement would still be better off dedicating their time to writing traditional articles.

Having said this, it is noteworthy that publishing in prestigious venues (acquiring "prestige points") is not the only conceivable route to professional success. Thanks to the Internet, everyone can access any article at any time: publication in journals is no longer motivated by the need to transport one's words to potential readers. Instead, academics seek publications in prestigious venues in order to catch the attention of potential readers and/or attain the rewards which a strong curriculum vitae tends to bring. Another way of achieving either objective is through direct reputation-building (without the intermediary of prestige points). One can in principle have an excellent reputation as an academic without publishing in any prestigious venue, and one could in principle attain the same rewards in this way. The Saul Kripke of 1971 (before the publication of *Naming and Necessity*, but after the Princeton lectures) nearly exemplified this: while Kripke had a number of publications in the prestigious *Journal of Symbolic Logic* by 1971, his reputation went far beyond their cumulative prestige; hundreds of individuals had scored comparable numbers of prestige points without attaining a fraction of his reputation.

However, it is doubtful that Kripke's career would have taken off like it did without his precocious publications in the prestigious *Journal of Symbolic Logic*. There are two main reasons why prestige points are in practice essential to professional success for an academic (even for Kripke). First and foremost, they are the easiest, fastest means of gaining reputation. For the purposes of reputation-building, a peer-reviewed publication functions like an eminent academic willing to read one's work on demand and broadcast his or her opinion of it to the world at large *if sufficiently positive* (and without the risk of annoying the recipients of the message). Word-of-mouth is a much slower, riskier means of building a good reputation. The result is that, in practice, the only way to stand out is to acquire considerable prestige points. Secondly, resources are often allocated in good part on the basis of prestige points independently of reputation. This is clearest in the case of large funding bodies, which tend to give considerable weight to publications in their point systems for grant allocation. These points are attributed quite independently of other considerations. And resources gained in this way typically come with their own prestige points in a positive feedback loop. These factors combine to make publications in prestigious venues the only viable route to gaining a readership and achieving professional success.

In summary, prestige points play a central role in today's academic ecosystem because they provide the only satisfactorily reliable, simple, content-agnostic, and general means of ranking and assessing authors and their contributions for the purposes of allocating attention and other resources. Simple, content-agnostic, reasonably reliable metrics such as prestige points are critical in guiding researchers, funding bodies, and all other decision makers in the academic ecosystem in their allocation of attention, trust, money, and other resources. For this reason, it seems likely that online discussions would have to generate prestige points or lend themselves to similar

forms of quantitative assessment in order to play a major role in communications between academics in philosophy.

2   Online surveys: potential and reality

Experimental philosophers have used surveys of the intuitive judgments of laypersons to defend certain philosophical views. Here I am interested in another kind of survey: surveys aimed at establishing how certain philosophical views are distributed among philosophers. An example of a survey of this kind is that David Chalmers and I recently conducted: The PhilPapers Philosophical Survey. The official invitation sent to some 2000 philosophers read as follows.

> The aim of the Philosophical Survey is to discover information about the distribution of views in the profession. Philosophers often make claims such as "Most philosophers believe P", "Most philosophers who believe P believe Q", and "Most philosophers of type T believe P". We aim to provide data regarding claims of this sort. [...]

> We will ask thirty questions, each giving a choice between 2-4 views on a philosophical issue (for example, "Analytic-synthetic distinction: yes or no?"; "Mind: physicalism or non-physicalism?"). Respondents can indicate that they accept or lean toward one of the options or can give one of a variety of "other" answers. [...] We will also ask for some optional background information.

Figure 1 shows one survey question as it appeared on the Survey site. The options

Figure 1: Question as presented on the Survey site

available under "other" for the illustrated question were:

- Accept both

- Reject both

- Insufficiently familiar with the issue

- The question is too unclear to answer

- Agnostic/undecided

- There is no fact of the matter

- Accept another alternative

- Accept an intermediate view

- Other

Slightly different options were available for questions which presented respondents with more than two main options. Respondents were also allowed to leave questions unanswered. The order of questions was randomized. So was the order in which answer options were presented.

The PhilPapers Philosophical Survey was open to anyone to take. However, we actively sought the participation of certain sub-populations in order to insure maximally informative results for these sub-populations. In particular, we defined a primary target population consisting of the 1974 faculty members of 99 leading analytic departments.[2] The participation of each of these 1974 individuals was solicited by direct email. 931 members of this group completed the survey (3226 individuals from all populations completed the survey). Assuming there is no sampling bias, we can estimate the distribution of views among the primary target population within maximum 2.4 percentage points for all views (often much less; confidence 95%). Please see the Survey official report[3] for more details on the Survey's questions and methodology.

It is worth emphasizing that the PhilPapers Survey was not aimed at determining the correct answers to the questions asked. The aim was merely to determine the distribution of views on the questions asked. We believe that this sociological information has an important role to play even if answers to philosophical questions cannot be inferred from it.

The reason we believe this information is valuable is that philosophers' beliefs about the views others hold inevitably play an important role in determining what topics they will discuss and what they will say about them. For example, if a theory already makes consensus, one's time might be better used than seeking to further the case for it (this may not always be the case, but this is often the case). Also, whether one is writing a journal article, a book, or a forum post, one almost always has to leave some premises undefended. One therefore needs to know to what degree the premises one plans to rely on are accepted by relevant members of the community; otherwise one might end up putting forward an ineffectual argument. These are two obvious ways in which beliefs about who holds what positions inevitably shape philosophers' public contributions. Insofar as such sociological beliefs shape the field, we would certainly be better off with accurate sociological beliefs.

---

[2]The departments chosen were those above a certain threshold in the Philosophical Gourmet Report of English-speaking universities, plus a handful of leading analytic departments outside the English-speaking world. See the Survey report (link below).

[3]http://philpapers.org/surveys/

As things stand, it is very hard to form a detailed, accurate view of philosophers' positions in debates one is interested in. In general, one does not broadcast one's views. One's views on any given topic are only widely communicated (and in a very limited way) when one publishes a paper in which they play an explicit role. And few broadcast retractions after having publicly endorsed a certain view. The result is that the vast majority of philosophers must base their assessments of their peers' positions on very small, potentially highly biased samples of the overall population (those people they have had the chance to talk to directly) combined with a distorted, faint echo of what goes on in the minds of philosophers worldwide found in published articles. We are by and large in the dark regarding what our peers believe.

This assessment is born out by the results of the PhilPapers Metasurvey we conducted as a follow-up to the Survey. The Metasurvey asked respondents to predict the results of the Survey. Participants had to give their best estimate of the distribution of views on the thirty central questions among the primary target population (the faculty members of the 99 departments mentioned earlier). 216 members of the primary target population who took the Survey have also completed the Metasurvey (as well as 511 others, but I will focus on the responses of the first group here). I will examine the Metasurvey performance data for the Survey question on the analytic-synthetic distinction as an example. Data on a few other questions which show particularly interesting response patterns are found in table 1.

The original question on the analytic-synthetic question was formulated as follows in the Survey: "Analytic-synthetic distinction: yes or no?". Metasurvey respondents had to estimate the proportions of individuals in the primary target group who would choose answers in each of the three following sets:

1. Accept: yes; Learn toward: yes

2. Accept: no; Learn toward: no

3. An option under "other"; no answer (skipped the question)

So respondents to the Metasurvey had to specify three percentages: one for "yes" views, one for "no" views, and one for "other" views.[4]

Many respondents reported difficulty assigning a percentage to the "other" positions, and this turned out to be one of the main causes of incorrect estimation (in general, respondents significantly underestimated the percentage of "other" answers). Since this cannot be regarded entirely as an indication of a shortcoming in knowledge of others' beliefs (some "other" answers reflect an unwillingness to answer rather than a philosophical belief), it is best to consider only the estimated and actual ratios of "yes" and "no" answers. To this end, we can normalize estimated and actual "yes" and "no" percentages so they sum to 100%.

Let an individual's *error* for a given view be the difference between their normalized estimated percentage for this view and the normalized actual percentage for this view. The average error for the "yes" view on the analytic-synthetic question is -20.6%. Since this question offered only two main options (which was not true of all

---

[4]Respondents were required to make the three percentages sum to 100%.

Table 1: A sample of Metasurvey errors (normalized; figures rounded to one decimal)

Question	Answer	mean est.	actual	mean err.	|err.|>20%	|err.|>30%
A/S distinction	yes	50.0%	70.6%	-20.6%	58%	35%
(same)	no	50.0%	29.4%	20.6%	(same)	(same)
Aesthetic value	subjective	67.7%	45.7%	22.0%	63%	34%
(same)	objective	32.3%	54.3%	-22.0%	(same)	(same)
Moral judgment	cognitivism	60.1%	79.5%	-19.4%	49%	22%
(same)	non-cognitivism	39.9%	20.5%	19.4%	(same)	(same)
Knowledge claims	contextualism	43.4%	42.0%	1.3%	25%	8%
(same)	invariantism	39.0%	54.1%	-15.0%	43%	12%
(same)	relativism	17.6%	3.9%	13.7%	23%	8%

questions), the average error for the alternate main option ("no") is the opposite, i.e. 20.6%.[5] The average individual in the comparatively well-informed primary target group expects a 50/50 normalized distribution on this question, but in fact the normalized percentages for the "yes" and "no" views are 70.6% and 29.4%, respectively.

Perhaps even more telling than average error is the number of respondents with high absolute errors. Percentages of respondents with absolute errors above 20% and 30% are given in the last two columns of table 1. These high percentages are significant. In the case of the analytic-synthetic distinction, an error of -20% implies that the respondent assigned a 50.6/49.4 distribution instead of the correct 70.6/29.4. This is the difference between being in a position to assume the analytic-synthetic distinction in many contexts and not being in a position to assume the analytic-synthetic distinction in most contexts. A majority of respondents from what are arguably the 99 most reputed analytic departments in the world had an absolute error of at least 20% for this question. An error of -30% implies that the participant takes the analytic-synthetic distinction to be rejected by a comfortable majority of respondents (59.4%) in the primary target group, while in fact it is accepted by 70.6%. 35% of respondents from the same group had an absolute error of 30% or more.

While there seems to be a need for more sociological data, surveying philosophers on their take on well-know "isms" only goes so far. A number of participants to the Survey have pointed out that its dichotomies left out important positions. Many who chose to endorse views offered have remarked that the options they chose did not reflect the subtlety of their actual views. Many others complained about the vagueness or unclarity of the positions as stated, which were generally expressed as isms without accompanying clarifications. We were fully aware of these shortcomings before launching the survey; we thought it would nonetheless be a valuable exercise. But there are serious challenges here for future surveys, especially if we hope to provide data on large numbers of specialized questions.

We can discern two related challenges for surveys. The first stems in part from the fact that surveys can only query individuals for their level of agreement on certain

---

[5]A t-test shows with confidence 95% that the mean absolute normalized population error for the primary target group (1974 individuals) on this question is at least 18.3% ($s = 19.6$; N= 201).

views without elaboration on their actual positions (otherwise it would not be possible to compile statistics). This has for consequence that the questions must identify positions with which a vast majority of individuals either agree or disagree. If the questions do not cut the debate along the right lines, answers will be uninformative (respondents will have to abstain from agreeing or disagreeing with the views proposed). But it is hard to know how to cut the debate. The second challenge is to offer sufficiently unambiguous, precise questions so that it is possible to draw illuminating conclusions about our peers' beliefs based on their answers—it must be possible to tell what *thought* someone who agrees with a certain statement actually agrees with. This, too, is difficult, because one of the principal challenges of academic philosophy is that the isms and other widely discussed theses are generally vague or ambiguous in important respects. Uncovering these formulation problems is a big part of what one has to do as a philosopher.

The PhilPapers Survey was meant in part as a prototype for more comprehensive, more probing surveys in all areas of philosophy. One thing it has taught us is that we need to find some way to systematically handle the preceding difficulties to make large-scale surveys possible.

3   The philosophical registry

In the previous sections I have made a case for a greater place for online discussion spaces and surveys in professional communications among philosophers, but I have also recognized major obstacles to the widespread use of either mode of communication. In this section I want to sketch a possible communication system which integrates the two in a way that might help overcome the difficulties I have highlighted: the philosophical registry.

The philosophical registry would contain three main kinds of object:

Propositions Propositions are items one can hold to be true or false.

Definitions Definitions serve to specify propositions.

Views Views record participants' take on propositions.

Propositions would come in two varieties: basic propositions and metapropositions. For now I will only talk about basic propositions.

A basic proposition would be a relatively brief natural language statement, e.g. "all facts supervene on physical facts". Each basic proposition would optionally have associated definitions which would be partly definitive of the proposition (orthographically identical propositions associated with different definitions would count as distinct propositions). Definitions could themselves be associated with further definitions in order to clarify expressions occurring in their definiens. However, definitions would not be allowed for theory names (e.g. "physicalism"). Instead, participants would be encouraged to submit distinct propositions for distinct definitions of the theories they are interested in. For example, someone interested in the thesis that all facts supervene on physical facts (arguably one reasonable understanding of "physicalism") would be encouraged to submit this claim as a basic proposition instead

of submitting a proposition which says "physicalism is true" and defining "physicalism" as stating that all facts supervene on physical facts. As many variations on or interpretations of any given ism as necessary would be allowed in the registry.

Anyone would be able to propose propositions and associated definitions. However, at least initially, contributions would have to be vetted by members of topic-specific panels of referees before being made available for assessment. Referees would assess contributions for relevance, clarity and originality (in the sense that obvious duplication would be prevented whenever possible). Propositions would be assigned to general topics by submitters, and panels would be assigned based on the topics selected. Ideally, topics would be drawn from a comprehensive, standardized set.

The main use of the registry would be to record one's views by specifying which propositions one accepts and rejects. Rather than respondents picking among putative dichotomies of positions as in the PhilPapers Philosophical Survey, participants would express their views on propositions (basic or meta) individually. One would be able to take one of three stances on any proposition: "true", "false" or "neither". The first and second options would signify that one accepts or rejects the proposition, respectively. The third would indicate that one considers that the proposition can be given neither a "true" nor a "false" truth value, be it due to unclarity, ambiguity, a lack of relevant facts, presupposition failure, paradoxicality, or any other reason.

Unlike a traditional survey, which is open for a given amount of time only, it would be possible to record and change one's views at any time. Results would be updated frequently and historic data retained for comparison.

All participants would be able to record their views on every accepted proposition. However, it would be possible to separate results for different sub-populations. The following are examples of properties which could be used to distinguish populations depending on the information one needs from the registry: areas of specialization, publication profiles, education (e.g. having a degree in philosophy or not), region of education, employment (being an employed philosopher or not), place of employment (various lists of universities might be drawn upon), nationality, place of work, and age.

Most well known isms have multiple interpretations—sometimes, it seems, almost as many interpretations as participants to a debate. We might therefore expect the registry of basic propositions to grow rapidly given sufficient participation. This would make it necessary to organize propositions more finely than simply associating them with broad topics. This is the purpose of the metapropositions mentioned earlier.

Metapropositions would be propositions which encode relations between basic propositions. Their primary use would be to organize the space of basic propositions in a way that makes it as easy as possible for participants to find basic propositions of interest to them. At least initially, metapropositions asserting that one basic proposition is a variation on another might be enough for this purpose. The relation of variation could be defined as follows:

Variation for any two distinct propositions P and Q, P is a variation on Q just in case either P is a disambiguation of Q, P is a precisification of Q, or P has roughly

the same interesting applications and implications as Q.

The primary aim is for this relation to preserve interest: if one is interested in P and Q is a variation on P, one should in most cases also be interested in Q. Of course, distinct relations and metaproposition types could be defined for the three criteria which make up the above definition of a variation, but I suspect that it would in many cases be too difficult to decide which relations should be used.

Both basic and meta propositions would be available for voting upon. To the extent that there is widespread agreement as to which propositions are variations on which, it should be possible to structure basic propositions according to the variation relation when browsing or searching the registry. It should also be possible to use voting statistics to determine which variations on any given proposition are most likely to be of interest to the average registry user. In this way, it should be possible to let a large number of basic propositions into the system while at the same time enabling registry users to find the propositions which are of interest to them. In other words, the solution to the problem of cutting the pie "the right way" noted in section 2 might be to cut it all potentially relevant ways at once while empowering the community to sort out which ways are most useful and interesting.

So far, there is no role for discussion spaces. These would enter the system through justifications which participants would optionally attach to their views. Each attached justification would take the form of a reference to a content item. Initially, two types of item would be referenceable: contributions to online discussion forums or weblogs, and publications.

In section 1, I suggested that the principal obstacle to the widespread adoption of online discussion spaces as channels for research outputs is their incompatibility with current quantitative assessment methods. Services such as the philosophical registry could help surmount this obstacle.

The registry could help with the assessment problem by providing us with enough data about contributors to online discussions to enable quantitative, content-agnostic assessments which at least compare in reliability with current prestige-based assessment methods. Three kinds of information drawn from the registry might be particularly useful for this purpose:

Deferrals Who defers to who for justification can be an indicator of who provides the best evidence.

Attention Whose submitted propositions gather most interest can be an indicator of who contributes the most important content.

Agreement Whose registered views best predict consensus (or near-consensus, or convergence) in the long run can be an indicator of who is most apt to recognize or provide conclusive evidence.

Of course, all of these factors can reliably indicate the best contributors only to the extent that convergence in the voting and usage behaviour of the participants in terms of which they are defined is a good indicator of what is true or interesting. For this reason, not all votes and deferrals could be counted as equal. For example, a weblog article which has attracted 100 deferrals from unpublished individuals should be

worth less "epistemic credits" to its author than one which has attracted 100 deferrals from widely published academics. Initially, publication records and possibly other credentials would have to be factored into metrics derived from the above. Later, perhaps, such factors could be replaced by participants' existing credits as derived from the system (in a manner similar to the PageRank algorithm[6]).

There is no doubt that even the voting and usage behavior of the most published philosophers does not always indicate the true or the interesting. Any quantitative assessment method based on the above would have to be used cautiously. But note that, as far as making a content-agnostic, quantitative assessment of the worth of someone's work goes, one has no choice but to rely on consensus or convergence among a certain population of trusted participants as a guide to the true or the interesting. If one cannot use the philosophical judgment of anyone in particular (including oneself), an assessment based on the voting and deferral behaviour of a trusted population over time is the best that one can do. Insofar as content-agnostic quantitative assessments must play a role in academia, assessment methods based on the above factors might be the best we can hope for. Note in particular that such methods could easily be more reliable at identifying good contributors and contributions than the current system of prestige points, especially when assessing junior academics with relatively few contributions to their record.

Once again, I do not believe that deferrals, voting behaviour, publications counts, citation counts or any other simple metric can reliably indicate who is most deserving of any given resource (attention, a job, a grant, etc.), but I recognize that such metrics in fact play an important role in determining resource allocation among academics, and that a professional communication system has to support such assessment methods to be sustainable.

It goes without saying that metrics derived from a service such as the philosophical registry would be of little interest until such a service became widely used by the community. For this reason, the registry is not a solution to the quantitative assessment problem for online discussions in the short to medium term. However, there is a chance that the added visibility a service such as the registry could bring to online discussions could stimulate contributions enough to take them to a level where registry-based metrics begin to see some use alongside traditional quantitative assessment methods. From that point on, online discussions and the registry could conceivably grow together at an accelerating pace in a positive feedback loop.

4 Conclusion

Books, articles, and talks play a central role in professional communications among philosophers. I have argued that other means of communication also ought to play important roles they do not play today: online discussion spaces and surveys, broadly conceived. I have tried to identify the main challenges to a widespread use of either type of service: online discussion spaces do not support current (prestige-based)

---

[6]Page, Lawrence; Brin, Sergeyl Motwani, Rajeev and Winograd, Terry (1999). "The PageRank Citation Ranking: Bringing Order to the Web." Technical Report. Stanford InfoLab. http://ilpubs.stanford.edu:8090/422/

quantitative assessment methods, and effective online surveys are extremely difficult to design, especially if one hopes to cover a wide range of views. I have proposed the philosophical registry as a community-driven service that could integrate a survey system with discussion spaces in a way that could help overcome these obstacles.

It is noteworthy that a registry of the kind sketched here could be particularly well suited to the application of a certain analytic methodology to philosophical problems. In some circles, it is widely believed that philosophical disputes are fueled in large part by linguistic confusion—either the participants have a poor grasp of the meanings of the words they use, or they have a poor grasp of others' meanings. On this view, the primary task of the philosopher is to dissipate the fog of language around the ideas under dispute. For most philosophical disputes, the closer we come to accomplishing this, the closer we will come to consensus. The philosophical registry could be particularly congenial to philosophers of this methodological persuasion. If analytic philosophers of the kind I am referring to are right, we might expect the simplest basic propositions on offer in the registry (named theories such as "physicalism", "incompatibilism", etc) to end up flagged as neither true nor false by large numbers of participants. They would then be superseded by variations which remedy their defects in meaning. This process would be iterated until (in many cases at least) the remaining propositions are all either obviously true, obviously false, or clear enough for empirical tests to settle. A registry such as that described here would not necessarily be inhospitable to those who reject this methodology, but it could conceivably serve as a catalyst for the efforts of those who appreciate its virtues.