What's in a (Missing) Name? Status, Quality and Attention in Open Standards Development *

Timothy S. Simcoe
J.L. Rotman School of Management, University of Toronto

Dave M. Waguespack
Robert H. Smith School of Business, University of Maryland

February 2009

*We gratefully acknowledge financial support from the NET Institute (www.netinst.com) and the Kauffman Foundation. Lee Fleming was a catalyst for this research and provided valuable input on several early drafts. We also thank Ajay Agrawal, Scott Bradner, and Brent Goldfarb for helpful comments. Address for correspondence: Joseph L. Rotman School of Management, 105 St. George Street, Toronto, ON M5S 3E6,Canada. E-mail: timothy.simcoe@rotman.utoronto.ca

What's in a (Missing) Name? Status, Quality and Attention in Open Standards Development

Abstract

How much are we influenced by an author's identity? If identity matters, is it because we have a "taste for status" or because it offers a useful shortcut — a signal that is correlated with the likely importance of their ideas? This paper presents evidence from a natural experiment that took place at the Internet Engineering Task Force (IETF) — a community of engineers and computer scientists who develop the protocols used to run the Internet. The results suggest that IETF participants use authors' identity as a signal or filter, paying more attention to proposals from high-status authors, and this has a surprisingly large impact on publication outcomes. There is little evidence of a "taste" for status. JEL Codes: L1, O3

1 Introduction

How much are we influenced by an author's identity as opposed to the actual quality of their work? If identity matters, is it because we have a "taste for status" or because it offers a useful shortcut — a signal that is correlated with the value of the work?

In a widely cited paper, Merton (1968) argued that reputation and identity are important factors in scientific careers. He provided qualitative evidence that famous scientists get a disproportionate share of the credit for joint discoveries, and labeled this phenomenon the "Matthew Effect" (after a Biblical reference to increasing returns).[1] One of Merton's anecdotes is based on a story about Lord Rayleigh, a British mathematician and physicist who won the 1904 Nobel Prize in Physics:

> "Rayleigh's name 'was either omitted or accidentally detached [from a manuscript] and the Committee turned it down as the work of one of those curious persons called paradoxers. However, when the authorship was discovered, the paper was found to have merits after all."

This paper presents evidence of the Matthew Effect based on a large-scale replication of Rayleigh's experience. We focus on a natural experiment that took place at the Internet Engineering Task Force (IETF), a community of engineers and computer scientists who develop the protocols used to run the Internet.

Between 2000 and 2003, over 5,900 papers were submitted to the IETF's "open" review process. For each submission, the IETF community received an email announcement providing basic bibliographic information such as the abstract and author names, and providing a hypertext link to the full paper. We exploit the fact that many authors' names were replaced by "et al" in that initial email announcement. Critically, the use of "et al" was variable and occurred solely at the discretion of the IETF's clerical staff, and not the direction of authors or any administrative policy (Fuller, 2006). This created a situation where some authors' names were randomly removed from announcements, while others' were left in place.

For submissions that have a low *ex ante* probability of success, there is a statistically significant drop in the probability of eventual publication when "et al" obscures the name of an IETF Working Group chair (i.e. a high-status author, roughly equivalent to a journal editor). This result provides compelling evidence that, at least in this case, the effects of status on screening are not driven by unobserved quality differences. Moreover, the magnitude of this

[1] "For unto every one that hath shall be given, and he shall have abundance: but from he that hath not shall be taken away even that which he hath." (Matthew 25:29)

result is striking, given that "et al" only removes authors' names from an email announcement; it does not result in a blind review process.

To address our second question — does this reflect a "taste" for status, or an informational shortcut? — we examine citations to published submissions, which are called RFCs. If the screening process discriminates in favor of high-status authors, their RFCs should receive fewer cites, since even their poor proposals are likely to be accepted. However, we find that status has no effect on RFC citation rates from patents or academic journal articles, while high status authors receive significantly *more* cites from other (future) RFCs. This suggests that while status is used primarily as a signal in the screening process, there is some taste-based favoritism in the IETF citation process.

We link these two results by proposing that attention is an important mediating variable that changes the size and nature of any status effects. In the screening process, status is used as a quality signal because IETF participants are faced with a flood of new submissions: on average 80 new or revised technical manuscripts, averaging 23 pages apiece, were submitted to the community every week. At the same time, high-status authors receive more citations (though only within the IETF community) because status provides an efficient mechanism for the IETF community to coordinate its actions. To further support the conclusion that author status influences the allocation of attention we return to the full sample of documents proposed for publication, and find that documents with unlisted high status authors are discussed less frequently on the IETF's electronic forums.

1.1 Status: Theory and Evidence

As the Rayleigh example illustrates, in situations where quality is uncertain, parties will frequently look to other social cues when making screening decisions. The social cue Merton focused on is position in the hierarchy of academic prestige. The selection committee cannot know for certain how Rayleigh's paper will ultimately be received by the larger scientific community, but can reasonably infer that work from a famous scientist will likely receive more attention and accolades. The critical point is that while position in the prestige hierarchy is somewhat driven by quality — Lord Rayleigh was undoubtedly a talented scientist whose work was influential — being well-regarded has its own rewards. Indeed, a central tenet of sociology is that individuals derive rewards from their social position, and not just their individual attributes (Merton, 1968; Gould, 2002).

To reiterate, the presumed distinction between object and social signal hinges on the notion that object attributes are frequently difficult or even impossible to ascertain prior to consumption. An employer can only estimate future on the job performance and trustworthiness of a potential new hire. A shopper can only forecast the future reliability of a new car. Given *ex ante* uncertainty, economists have famously noted noted two reputation related phenomenon that are surprising in light of conventional expectations about market-based transactions. First, for consumers of goods, dealing repeatedly with a reduced set of known reputable parties instead of engaging in search becomes a reasonable strategy for mitigating quality uncertainty and opportunism (Akerlof, 1970). Second, for high quality organizations and individuals producing goods or services, the difficulty potential customers have in assessing their offerings makes it reasonable for the producer to undertake costly actions that are not directly related to competitive goals in order to build a positive reputation (Spence, 1973). In both cases the reputational signal is economically useful because it is correlated with underlying merit.

Recent work coming from a sociological perspective has focused on quality signals derived from observable relationships, often explicitly defining status as position in a social network. Status makes the most intuitive sense as a construct distinct from reputation in those situations where the reputational signal is weak or non-existent, such as when an actor has no track record, or the consumer and producer have had no direct contact, or the producer is moving into a new domain. In these cases third parties make inferences about unobserved quality based on how well connected an actor is. Even if status predicts average unobserved quality accurately, such inferences are likely to be quite noisy as the observer is now two steps removed from the object. As such the notion that there is an "unearned" component of the quality signal is an important part of theorizing about status (Washington & Zajac, 2005). Podolny (1993) argues that status represents "perceived quality", where the status signal and actual quality are loosely linked, particularly over time. Bothner et al. (2009) posit that observers cannot confidently assume a relationship between status and quality, particularly if status based deference generates complacency on the part of the producer.

Further, from the status perspective, there are three reasons for concern that seemingly useful quality signals will morph into relatively rigid self reinforcing social hierarchies. First, biased screening can make it unlikely that some alternatives are ever considered (Podolny, 1993; Podolny et al., 1996; Burris, 2004). In essence, and even while in theory reputations will rise and fall, if the bias for status is sufficiently strong then consumers may simply never observe the counterfactual. Second, biased screening may influence the ex post expectations and behavior of the consumer (Sorenson & Waguespack, 2006). For instance, believing that a new employee from an elite university has greater potential, an employer may give her better training and opportunities would otherwise be the case. Even while the screening counterfactual is potentially observed, *ex ante* quality becomes conflated with disparate post transaction investments. Third, social structures can become categories that define a producer's own expectations about legitimate roles (Zuckerman, 1999; Phillips & Zuckerman, 2001). Merton (1968) notes that what sets prominent scientists apart is not just the potentially biased attributions that others make to them, it also their own propensity to confidently take on new and important research problems. If lower status actors self-select out of more valuable activities, the relationship between ex ante quality and performance again becomes complicated.

In sum, a growing body of recent work has demonstrated that status functions as a form of positive discrimination, generating increasing returns to the well positioned beyond individual merit (Stuart et al., 1999; Benjamin & Podolny, 1999; Washington & Zajac, 2005). Academic work on status and related social relationship dynamics has blossomed — Podolny's 1993 article alone has received over 280 citations. In a review of social capital research, however, Mouw (2006) notes that even while theoretical distinctions have become increasingly refined, the field lacks the data to adjudicate among those alternatives much less solid evidence that there is a real *causal* relationship between social capital of any form and outcomes. We posit that finding evidence of status effects, and hence developing the theoretical literature, is hampered by two substantial empirical challenges.

First, status, or any other measure of third party regard, is empirically difficult to separate from actual *ex ante* quality. The typical approach in archival studies is to test for status effects while throwing in as many controls as possible for quality. As status is not experimentally assigned in these archival studies one cannot rule out the alternative that status effects are simply driven by unobserved (to the researcher) *ex ante* quality differences. Another way of stating the problem is that the causal arrow between status and performance may run in the opposite direction, with variation on status simply an outcome of variation on performance, and not the other way around (Gould, 2002). One can imagine that the unobserved quality concern is exacerbated in highly complex situations, precisely the situations in which status is hypothesized as most relevant, because in these settings the relative comprehension of the academic researcher is most diminished. In other words, in complex settings the academic researcher is himself a layperson searching for social cues that are correlated with *ex ante* quality. It remains unknown whether true experts themselves rely on these social cues, or use other correlated indicators that were not included in the academic models.

Second, it is difficult to draw conclusions about mechanisms, whether status effects reflect useful quality signals or are driven by tastes, without examining both short run screening outcomes – i.e. publish or reject a paper? – and longer run *ex post* outcomes – i.e., how well cited is the published paper? In the economics literature on disparate treatment the interpretation is that taste-based discrimination results in suboptimal *ex post* outcomes because preferred actors face a lower *ex ante* selection threshold (e.g. Becker, 1993). Elites, over-confident in their own abilities and therefore less likely to respond to external changes or develop new skills, may ultimately have reduced performance (Bothner et al., 2009). Conversely, if status leads to better or no worse *ex post* outcomes, then this provides some evidence that status signals are useful search shortcuts. In the quality signal interpretation screeners do not intrinsically prefer high status parties, its just that exhaustively searching all of the available alternatives becomes very expensive and may yield no better results *ex post*. The quality signal interpretation holds even if *ex post* quality is conceived of as at least partially endogenous, consisting of socially influenced judgements about relative merit rather than some external truth, in which case status hierarchies are essentially mechanisms for social coordination (Gould, 2002).

In this paper we heed the call of Manski (2000) for better data and clear thinking, by developing clear empirical tests for relatively simple status constructs. To address the causality concern we employ a natural experiment — the suppression of high status author names on working papers. If author status matters for publication decisions, then high status author name suppression will make publication less likely. Assuming "et al" use is random then omitted variable endogeneity is not a concern, and we have persuasive evidence for a status effect. To address the issue of mechanisms, we examine the short run publication outcome, the short run outcome of the amount of email attention generated, and longer run future citations within and outside the focal community. If author status serves to coordinate attention, then high status author name suppression will diminish the amount of discussion about a paper. If author status is a useful quality signal, then *ex post* citations for high status authors will be greater or no worse.

This paper is related to earlier work on discrimination in the academic publishing process. For example, Blank (1991) conducted a randomized experiment on the impact of double-blind refereeing at the American Economic Review. She found that double-blind procedures led to true anonymity for only 54 percent of the referees, and increased the acceptance rate for authors from the most highly ranked institutions. Smart & Waldfogel (1996) also found evidence of reverse discrimination by comparing published papers' citations to their observed editorial treatment (e.g. page length and placement within the journal). Peters & Ceci (1982) conducted a field experiment by re-submitting twelve papers published by "prestigious" authors to the psychology journals where they were originally published. Nine of the re-submissions went undetected, and of those eight were rejected on re-submission.

Our work also fits into the broader literature on discrimination. This paper is one of a handful that combine data on selection and performance. Other examples include Ayres & Waldfogel (1994) on discrimination in bail-setting, and Hellerstein & Neumark (2004) on labor-market discrimination. Unlike many papers that that focus on the selection process, we identify a natural experiment — the random deletion of names by "et al" in announcements — that allows us to get around the problem of omitted variables. Randomization of names has also been used to study racial discrimination in field experiments, such as Bertrand & Mullainathan

(2004). Finally, the fact that "et al" has a large impact on the chance of publication — even though it is not particularly hard to find a complete list of authors — makes our paper relevant to the behavioral literature on "salience" or attention costs (see DellaVigna, 2008, for an overview).

The remainder of the paper proceeds as follows: Section 2 describes the IETF publication process and the "et al" natural experiment. Section 3 describes our data, measures and statistical methods. Section 4 presents results. Section 5 concludes.

2 Setting and Experiment

This section describes the IETF and the process that led to exogenous removal of author names from new proposal announcements.

2.1 The IETF Publication Process

The IETF is a non-profit organization that creates compatibility standards used to run the Internet. These standards provide a set of rules for new product design. By adhering to IETF standards, engineers can ensure that new products will be able to acquire an Internet address, encrypt a message, display a web page, play a video, or identify an e-mail attachment.

Anyone who is interested can become an IETF member simply by choosing to participate, though most members are engineers or computer scientists. While the IETF holds three plenary meetings each year, most work takes place on the Internet, where new ideas are introduced and debated on a series of email listservs. These listservs are often administered by an IETF Working Group (WG) that is developing new standards for a particular technology area.

Figure 1 provides an overview of the IETF publication process. Anyone can submit a proposal. These submissions are called Internet Drafts (IDs). An ID usually goes through a series of revisions before it is either discarded or accepted for publication. Published drafts are referred to as Request for Comments (RFC).

There are two ways for a draft to become an RFC. The first is to become the "consensus" recommendation of an IETF Working Group. This implies support from a solid majority of WG participants (though not unanimity). In practice, the Working Group chair usually determines whether there is a consensus. If a chair decides there is consensus support for an ID, it is forwarded to the Internet Engineering Steering Group (IESG) — a sort of editorial review board. The IESG then issues a last call to the entire IETF. If the last call raises serious issues, the Internet Draft is sent back to the Working Group. Otherwise it becomes an RFC.

The second way for a draft to become an RFC is through the "independent submission" process. In this case, an individual posts a new Internet Draft to a public web server maintained by the IETF. If the draft fits into the agenda of an existing Working Group, or there is sufficient interest in starting a new Working Group, the proposal might become a Working Group draft (though this rarely happens in practice). If the draft does not get picked up by a Working Group, but there is enough interest within the IETF community, the individual may submit it to the RFC editor and ask that it be published as an individual (non Working Group) RFC. The RFC editor typically sends these drafts out to subject matter experts for review. If the reviewers and RFC editor agree that the ID should be published, it is sent to the IESG to ensure that it does not conflict with the work of any IETF Working Groups. If the IESG approves, the draft is published as an RFC.

The Working Group process is used to publish more important RFCs. For example, only WG drafts can be formally designated as IETF Standards. Individual submissions are classified as either Experimental or Informational RFCs. Many WG drafts are essentially "commissioned" works, and when an individual submission is recognized as particularly important, it will generally be adopted by a Working Group. Because WG drafts are more likely to be commercially significant, the approval process can be political.

Still, RFCs published via the individual submission process can be influential. They often propose new uses for IETF technology, describe lessons learned from implementation, and pose new problems for IETF members to work on. The IETF receives a large number of unsolicited IDs, and RFCs published via the individual submission process are cited frequently. We focus on individual submissions in the statistical work below.

Internet Drafts that have not been revised in more than 6 months are said to have "expired" and are removed from the web page where all current IDs are posted for review. While it seems clear that some drafts on both the WG and individual-submission track are essentially rejected (i.e. informed the revision is a waste of time) we can only observe expiration — which might reflect either outright rejection or a decision by the author's to abandon the project.

2.2 The "et al" Experiment

When a new Internet Draft is submitted to the IETF, it is posted on a web page where anyone can download it. With the initial posting (and for each subsequent revision) an e-mail announcement is sent to the entire IETF via the "ietf-announce" listserv. The top panel in Figure 2 shows a typical message sent to the ietf-announce list. The message contains a title, list of authors, filename (which contains information about WG affiliation and revision history), date, and abstract. The bottom panel in that figure shows that the same information is available in the ID text. In particular, the authors' name and institutional affiliations can be found in the header; while detailed information is typically available at the back of the document.

Before 1999 every announcement included the name of every author who contributed to the proposal. However, beginning in 2000, there was a rapid increase in the use of "et al" in the author list when IDs had more than one author. Our identification strategy exploits the fact that the names obscured by "et al" were occasionally prominent members of the IETF community. Figure 3 shows an ietf-announce message where three author names were replaced by "et al" and one of the names, which appears in the document header, is a former Working Group chair. We also discovered a number of cases where more authors were listed on the ID than the announcement, but there was no "et al" in the announcement to indicate this fact. We refer to these cases as "invisible et als."

The decision to use "et al" was made by the IETF Secretariat — an administrative body that manages the ietf-announce listserv. The secretariat had a small staff that would process incoming drafts by typing the relevant information into a standardized form. Conversations with the director of the IETF Secretariat (Fuller, 2006) suggest that the decision to begin using "et al" for co-authored IDs was simply a response to the rapidly growing volume of submissions during that time period. Figure 4 shows the dramatic increase in ID submissions — particularly individual submissions which more than doubled between 1998 and 2000.

The decision whether to use "et al" on a particular draft was left up to individual clerical staff who typed in the ietf-announce messages. These individuals suggested that they tried to include every name. However, they would resort to "et al" when things became extremely busy. This typically occurred during the time-period immediately before an IETF meeting, when they often received a flood of new proposals. Figure 5 provides strong evidence to support these claims.

The top panel in Figure 5 is based on a kernel-density regression of "et al" usage on time (the unit of observation is a new ID submission with more than one author). The vertical bars represent IETF meeting dates. It is obvious that there is both an increasing trend towards using "et al" and a strong cyclical component that is tied to the meeting dates. The second panel in Figure 5 overlays the predicted "et al" probabilities with the results of a second kernel-density regression showing the rate of new submissions over time. From this figure, it is quite clear that submission volume is also driven by the existence of meeting-related deadlines. However, it is also clear that between late 2000 and mid-2004, any proposal with 2 or more authors had a 30 to 60 percent chance of receiving an "etal" in the announcement process.

We study the impact of individual status on publication outcomes by assuming that the use of "et al" produces an exogenous shock to the information set of ietf-announce readers. In particular, we study what happens when a "high status" name is removed from the message announcing a new draft. Note that this is an extremely weak treatment. In particular, any interested reader can obtain a complete list of authors by simply downloading a copy of the proposal and glancing at the front page. In most cases, the ietf-announce message contains a link to the draft, so this is very easy to do.

Our primary measure of status is a dummy variable that indicates whether an author has ever chaired an IETF Working Group. The position of WG chair is similar to journal editor in many respects. Chairs decide whether and when a Working Group has reached consensus (though this decision is subject to IESG review). Consequently, these individuals have high visibility within the IETF.[2] Like journal editors, chairs are usually awarded their position on the basis of past publication success. In particular, Fleming & Waguespack (2009) show that past RFC publications increase the likelihood of becoming a chair.

3 Data and Methods

This section provides an overview of the data set and discusses the empirical methods used below.

3.1 Data

Using public sources, we collected information on every Internet Draft submitted to the IETF between 1992 and 2004. This sample contains 12,342 Internet Drafts. For each proposal we have a submission date, a complete list of authors, and an outcome: expiration (failure) or publication as an RFC.

For each author, we have a complete list of the Internet Drafts they submitted to the IETF, their place in the list of draft authors, an email address, and an indicator of when and whether the author ever served as a Working Group chair. We use authors' email addresses to construct a number of additional variables. For example, we can observe whether a proposal has co-authors from two different institutions (e.g. cisco.com and Harvard.edu). Table 1 provides a complete list of variables and definitions.

Table 2 provides summary statistics for the estimation sample, which contains all individual IDs submitted between 2000 and 2003 with between two and five authors, and no WG chair in the lead author position.[3] These submissions had a very low 4.8 percent probability of becoming an RFC. On average, a new submission was discussed on 1.4 WG email listservs, and generated 4.3 email messages within 7 days of posting. Conditional on publication, an RFC is cited roughly 2.6 times by other RFCs, and 0.5 times by a US patent or academic journal.

[2] Chairs are not allowed to be an author on drafts submitted to their own Working Group.

[3] We drop submissions with 6 or more authors because the probability of "et al" approaches unity. There are relatively few such IDs, and they appear very different on a number of observable dimensions.

Roughly 26 percent of the estimation sample has one or more authors removed by "et al" while 8.2 percent of these IDs have an "invisible" missing name. However, only 6.8 percent of the observations have a missing WG Chair, resulting in a problem of small numbers. Since our empirical strategy is to compare cases with a Listed Chair to a Missing Chair, we typically use only a few hundred observations out of several thousand observations to identify the Matthew Effect. While our results suggest that the Matthew Effect is large, the estimates are often imprecise.

3.2 Methods

3.2.1 Identifying the Matthew Effect

We have data from a census of Internet Drafts (indexed by i). Our main dependent variable RFC_i is a dummy variable that equals one if draft i is ultimately published as an RFC. The variable $Chair_i$ is a dummy that equals one if Internet Draft i has one or more authors who are a current or former WG Chair. We construct a variable $MissChair_i$ that equals one for Internet Drafts where there is a chair author, but that author is not listed in the draft announcement. If we assume that "et al" is exogenous, a very simple statistical test for the Matthew Effect would consider the null hypothesis:

$$E[RFC|Chair = 1, MissChair = 1] = E[RFC|Chair = 1, MissChair = 0]$$

for drafts with the same number of authors where the "high status" Working Group chair author occupies the same position in the list of all authors. We begin with a non-parametric test based on this idea, for drafts where there are two to five authors, one of whom is a chair, where the chair is not listed first.

We then turn to a difference-in-differences model that estimates the impact of having a chair-author on the probability of publication, rather than conditioning it out. In particular, we estimate the following regression, where λ_t are a set of IETF meeting-date fixed effects (to capture time trends), $EtAl_i$ is a main effect, and X_i are a set of Internet Draft characteristics (including a constant):

$$RFC_i = MissChair_i\beta_1 + Chair_i\beta_2 + EtAl_i\beta_3 + \lambda_t + X_i\delta + \varepsilon_i \quad (1)$$

In this specification, β_2 measures the impact of having a high-status author on the proposal. The Matthew Effect is captured by β_1, which measures the change in publication rate when a chair's name is removed from the ID announcement (relative to the change when a non-chair's name is removed). Note that $MissChair_i$ is not quite an interaction between $EtAl$ and $Chair$;

in particular, it equals zero for some proposals where $EtAl_i = 1$ and $Chair_i = 1$ but it is possible to see the name of the Working Group chair because the "et al" cutoff is relatively high.

Most studies of status or discrimination run a model like (1) without the "et al" interaction and test the hypothesis that β_2 is zero. The standard critique of that approach is that it is not possible to correct for spurious correlation between $Chair_i$ and ε_i (e.g. because of omitted variables). In our model, we assume that "et al" administers an exogenous shock to the information set of IETF participants. If this shock is truly random, the "interaction term" $MissChair_i$ and the main effect $EtAl_i$ will be uncorrelated with ε_i. If IETF members treat "et al" as a low-status author, then our specification should produce an unbiased estimate of the Matthew Effect. If they form a guess by allowing that there is some probability that "et al" is a high-status author, we estimate a lower bound on the true effect.

3.3 Explaining the Matthew Effect

While the methods described above are an appealing way to identify the Matthew Effect, they are not particularly good at providing an explanation for it. Specifically, β_1 is an estimate of the the impact of changing IETF members' beliefs about proposal authors. To understand *why* these beliefs matter we would like to run a different experiment in which we randomly assign individuals to be Working Group chairs (or, less plausibly, hold the set of WG chairs fixed and randomly assign individual characteristics).

We can take one step in this direction by interacting $EtAl_i$ with author or proposal-level observable characteristics X_i, in a specification such as (1). This is akin to asking *which* authors or proposals produce the Matthew Effect. While this can provide useful evidence when there are a rich set of observable characteristics that a reader might use to discriminate between proposals, it immediately brings us back to the omitted variables problem — it is impossible to interact $EtAl_i$ with all of the potentially relevant author attributes. Put differently, even if we find a sizable Matthew Effect, it is not possible to discern whether IETF participants are screening on X or some (unobserved) alternative that is highly correlated with X.

An alternative approach to explaining the Matthew Effect is to search for evidence of the underlying mechanisms that lead to apparent discrimination. This seems particularly important in our case, where the treatment (omitting an authors name from an email announcement) is arguably quite weak. We suspect that a key mechanism in our setting is the amount of "attention" that a draft receives early in the submission and review process. A draft that receives very little attention seems more likely to be abandoned for a number of reasons: attention is needed for a Working Group to become interested in a proposal; authors may interpret lack of interest as an assessment of quality; and the quality of drafts that do receive early feedback may improve as a result. Given the rapid increase in total proposals, attention from prominent IETF members also seems likely to have been a scarce resource during our sample period.

We measure "attention" by counting the number of times an ID is mentioned in email messages to IETF listservs in the time-period immediately after it is announced. This measure focuses on replies to an initial message in order to distinguish a true conversation among IETF members from various automated messages that will mention a new draft. We use the difference in difference specification (1) to ask whether proposals where a high-status name is obscured by "et al" receive less attention than they otherwise would have.

A third approach to explaining the Matthew Effect is to use data on *ex post* outcomes to look for systematic differences across high and low status groups. The use of so-called "outcome tests" for discrimination was strongly advocated by Becker (1993), and Ayres (2002) provides an overview of the approach. In our setting, a natural measure of an ID's *ex post* performance is the number of citations it receives. Intuitively, if WG chairs get a "free pass" in the review process, the marginal chair-authored RFC should receive fewer citations than a non-chair authored RFC.

To be specific, suppose that cites are a valid measure of proposal quality and that the IETF's objective is to identify and publish high-quality proposals. The publication process produces a set of proposal-quality thresholds c_s, where $s \in H, L$ identifies an author's status. An unbiased screening process would generate a common set of quality thresholds $c_H = c_L$, so the distribution of cites conditional on publication is identical across status-groups. Taste-based discrimination in favor of high-status authors would lead to a lower quality-threshold for that group, so that

$$E[Cites|s = H] = E[Cites|q > c_H] \leq E[Cites|q > c_L] < E[Cites|s = L] \qquad (2)$$

While straightforward, the outcomes test has been criticized for several reasons. Most often, critics point to the fact that it makes predictions about the quality of a *marginal* as opposed to an *average* proposal. In order to operationalize the outcomes test, we must assume that proposals from high and low-status actors are drawn from the same underlying quality distribution. This assumption may not be very plausible. For example, if high-status authors produce work that is better on average, the outcomes test may show that average quality (conditional on acceptance) is equal even if there is discrimination against the low status types (i.e. $c_H < c_L$).[4]

[4]Knowles, Persico & Todd (2001) developed a formal model of discrimination in a search process which

A second, more obvious, criticism of outcomes tests (that has nevertheless received less attention) is that it is hard to determine whether an outcome measure is itself contaminated by taste-based discrimination.[5] One approach to this problem is to use prices as an *ex post* performance measure (as in Ayres & Waldfogel, 1994) — on the assumption that markets are unbiased (at least at the margin). Lacking an obvious market-based performance measure, we consider an alternative approach. Specifically, we compare citation-based performance measures from three different sources: future RFCs, US patents and academic publications. We would expect academic publications and especially patents to be less biased, since they are produced by individuals who are (on average) less involved in the IETF, and in the case of patents, have no incentive to draw attention to prior work.

However, our approach does raise questions about whether there is really a one-dimensional quality measure. In particular, the objective function that IETF members apply to the screening process may not coincide with the quality perceptions of academic paper-writers or patenters. Nevertheless, the use of multiple performance measures brings additional information to bear on the question of whether the Matthew Effect is statistical or taste-based when there is a potential for similar biases in the citation process.

4 Results

We begin by illustrating our main result using both non-parametric and parametric approaches. We then consider three extensions. First we interact the "et al" dummy with other variables to look for the specific factors that drive screening behavior. Then we examine how status influences attention as measured by emails. Finally, we conduct an outcomes test for taste-based discrimination using citations as a measure of publication quality.

4.1 The Matthew Effect

Table 3 presents non-parametric evidence of discrimination in favor of high-status Working Group chairs. We begin with a sample of all Internet Drafts submitted to the IETF between 2000 and 2003 with between two or five authors. We split this sample into individual IDs (top panel) and Working Group submissions (bottom panel). Note that publication probabilities are an order of magnitude larger for WG submissions. We find no evidence of a Matthew Effect

generates the prediction (2) even when there is unobserved heterogeneity — including differences in the quality distribution across groups. The key to their model (and the resulting "hit rate" test) is that high and low status authors respond to the incentives created by the search process when deciding whether to submit a proposal.

[5]For example, in the literature on discrimination in police searches, most authors use convictions as the outcome. Of course, this is problematic if the judicial system is prejudiced against certain defendants.

in this sample, and attribute that result to the fact that WG submissions receive a great deal of attention.

Table 3 also divides proposals into four columns: those with no WG Chair author; those with an unlisted chair-author; those with a listed chair-author; and those with a chair as the lead author. Our main result can be seen be reading across these column in the second row of Table 3: the publication rate of submissions with no chair author is very close to that of submissions with a missing chair (roughly 4 percent). However, if the chair-author is listed in the announcement, this probability increases to 14 percent, which is closer to the 19 percent success rate for IDs with a chair-author in the first position. (These means are different at the 1 percent level.) While the differences are smaller, the same general pattern holds for our email-based attention measures in the third and fourth row of Table 3.

We now turn to a parametric model of the Matthew Effect, based on equation (1). Table 4 presents these results. Since it is difficult to interpret interaction coefficients in a nonlinear model (Ai & Norton, 2003) we focus on linear probability models, though the second column also provides marginal effects from a logit regression (calculated at the sample means of X) for comparison.

The first column of Table 4 compares the second and third column in Table 3, controlling for the total number of authors and the submission date. We find that having a listed, as opposed to an unlisted, chair author increases the probability of RFC publication by 13 percent. This is a remarkably large estimate (especially given given the weak nature of the treatment effect) which suggests that revealing an author's identity can increase publication rates by a factor of three or four! We suspect that the Matthew Effect is particularly large in this particular setting, where a large number of proposals are being submitted, very few are likely to succeed and the rapid the growth of the Internet is placing demands on the attention of IETF participants.

The second column in Table 4 uses a logit specification to control for a number of observables, and finds that the Matthew Effect declines to 6.3 percent, but remains highly significant. The negative coefficient on Draft Queue suggests that publication rates decline when a large number of proposals are submitted in a short time window, which reinforces our notion that attention is an important factor in this setting.

One way to check that our estimates of the Matthew Effect are not spurious is to examine whether the use of "et al" also causes a large drop in the likelihood of publication when it does not obscure the name of a high-status author. The last four columns in Table 4 present results from a difference-in-differences model. Estimates of the Matthew Effect are somewhat more reasonable in this model. The main effect of having a high-status author is to increase the publication rate by roughly 8 percent. However, when a chair's name is not revealed, that increase is only half as large (between 3 to 4 percent).

One concern with our analysis is that we have not controlled for the order in which authors are listed, i.e. there might be a difference in (unobserved) quality if a chair is listed as the second instead of third author. We have removed all cases with a chair in the first position, since their name would always appear in the announcement. The last column in Table 4 provides an additional robustness check by restricting the sample to IDs with three authors. We find that our estimate of the Matthew Effect remains statistically significant.

4.2 Explaining the Matthew Effect

4.2.1 What's in a Name?

The previous sub-section showed statistically significant evidence of a fairly large Matthew Effect despite the fact that our "treatment" — removing a name from the email announcement — is rather weak. We now ask what might be causing this result? As a first step in this direction, we turn to the specification in equation (1), and interact the "et al" dummy with several variables that IETF participants might use to screen proposals.

The first column in Table 5 replicates our previous results on the Working Group Chair effect. The second column replaces the Working Group chair terms with a count of RFCs published by all authors on the draft, and an interaction term that counts RFCs published by authors who are not listed on the ietf-announce message. Doubling the number of RFCs published by the lead author increases the probability of becoming in RFC by 4 percent. There is a similar increase for RFCs published by non-lead authors who are visible on the announcement. However, the effect is only half as large (i.e. declines by 2.5 percent) when those RFCs are produced by unlisted authors.

The third column of Table 5 considers another piece of information that "et al" may be obscuring: the number of different organizations who sponsor the proposal. We find that drafts with authors from more than one institutional affiliation are 2.3 percent more likely to be published. However, when "et al" obscures the name of all but one organization, this effect disappears.

The last column in Table 5 adds all of these factors to a single model. The results suggest that most of the variation in publication rates is explained by past success. Indeed, becoming a WG Chair is highly correlated with publishing one or more influential RFCs. This is consistent with our "signaling" interpretation of these results: IETF members presumably understand that a high-status name is something that has been earned. To be clear, none of this information is available to readers of the ietf-announce messages, which only list the author's names. Thus, we should interpret these results as providing evidence on the average reader's knowledge of individual author attributes (based on their name) in addition to whatever quality inferences the reader draws based upon those attributes.

One of the things that is surprising about our findings is that readers appear to know a great deal about certain authors, and use this information to make informed guesses about proposal quality. Yet they do not bother to obtain complete author information (which is readily available in the actual proposals) even though they are generally "one click" away from learning this information. One interpretation of this finding is that in the early rounds of sorting through a large volume of papers that might be read, individuals have extremely high search costs. Our next set of regressions examine this idea more closely.

4.2.2 Status and Attention

Table 6 looks for evidence of a direct link between removing the name of a high-status author from the ietf-announce message, and the amount of attention that a proposal receives. We measure attention by counting the number of WG listservs where an ID is mentioned, and the number of times an ID is mentioned in a message, during a 7 day window after it is submitted. Our estimates are based on a negative binomial count model.[6]

Focusing on the first column in Table 6, we can see that there is a 32 percent increase in the number of listservs when a high-status author is listed on the ietf-announce message. However, two-thirds of that effect disappears when the chair's name is removed from the announcement. We also find that there is a large negative main effect of having "et al" on the ietf-announce message. This may reflect the fact that "et al" is used when there are a large number of new proposals arriving, so IETF members are less likely to pay attention to any individual new submission. The second column in Table 6 shows that the effect shrinks, but remains statistically significant, after controlling for observables including the presence of authors with past publication success.

The third and fourth columns of Table 6 show that there is a different pattern for Email replies. While having a high status author is associated with more attention, the effect does not shrink when that author is removed from the announce message. This may reflect the fact that WG chairs are active participants who send emails about their own IDS, whether or not they appear on the announcement.

The results in Table 6 suggesting that high-status authors get more attention, and a non-trivial piece of that attention is purely associated with having that author's name on the draft announcement. While this "attention effect" is plausible — particularly given the large volume and preliminary nature of many of these proposals — it is nevertheless surprising that a

[6]Hausman tests do not reject the hypothesis that the point estimates are the same as would be obtained from the less restrictive Poisson model. We use the negative binomial because it is efficient relative to Poisson.

difference of one or two email lists (or messages) leads to a substantial divergence in publication probabilities. We interpret these findings as evidence of strong increasing returns to attention in the early stages of the creative process. It is not clear whether this is driven by unique features of the IETF publication process, or is a more general feature of creative work. However, we believe it to be an interesting subject for future research.

4.2.3 Status and Citation

Our final approach to understanding the mechanisms behind the Matthew Effect is to conduct outcome tests. In particular, we estimate Negative Binomial models of RFC forward-citation rates to test the hypothesis that taste-based discrimination in favor of current or former Working Group chairs leads to a lower publication-quality threshold (and hence a lower forward-citation rate). We use citations from three different sources as a dependent variable: RFC cites, patent cites and cites from academic journal articles.

Table 7 presents the results. For RFC citations, having a chair author is associated with a significant *increase* in citations. This is not only inconsistent with discrimination in favor of these authors in the screening process; but suggests that there may be some favoritism at work in the citing behaviors. For the other two citation-types, the coefficient on a Working Group chair dummy variable is statistically insignificant. In part this reflect a smaller sample size; since relatively few individually submitted Internet Drafts are published as RFCs. However, only one of the point estimates is negative (and in that case, quite small).

The citation results suggests that the large "et al" effects we find above are driven by statistical rather than taste-based discrimination. In other words, IETF members are using names to screen out proposals (as the results in Tables 5 and 6 suggest) but not applying a different standard to the proposals submitted by different types of author.

5 Conclusion

Many authors have written about the importance of labels and identity. Perhaps the most famous statement of the hypothesis that a name does not (or should not) matter belongs to Shakespeare:

> What's in a name? that which we call a rose
> By any other name would smell as sweet;
> So Romeo would, were he not Romeo call'd,
> Retain that dear perfection which he owes
> Without that title.

This paper presents evidence that Juliet was wrong, at least within the context of Internet standards development. We exploit a unique natural experiment created by the fact that "et al" obscures the names of some authors who nevertheless contribute to proposals brought before the Internet Engineering Task Force. We find that when "et al" obscures the name of a high-status author — specifically a current or former IETF Working Group chair — there is a significant drop in the publication rate. When "et al" obscures the name of a low-status author, there is no change in the likelihood of publication.

These results provide statistical evidence of the Matthew Effect first described by Merton, and our natural experiment provides a unique opportunity to separate the role of identity from the quality of an author's ideas. Our estimates suggest that the Matthew Effect is surprisingly large in this setting, particularly given the relatively weak treatment — while "et al" obscures author names on an email announcement, it is relatively easy to find a complete list by downloading the relevant proposal.

We ask what gives rise to this large Matthew Effect; specifically, do author names serve as a signal of quality, or is there simply a "taste" for status that contributes to some individuals' success? Evidence from the "et al" experiment suggests that names are primarily used as a signal: the screening process is one of statistical rather than taste-based discrimination. This conclusion is based on two pieces of evidence. First, publication rates respond to other information that gets obscured by "et al" (e.g. whether how many RFCs have the authors published) and not just status cues. And second, citations to published RFCs suggest relatively little difference in the average quality of work produced by high and low status authors.

As an alternative to taste-based discrimination in favor of established authors, we suggest that increasing returns to attention may be a mechanism that explains how our relatively weak treatment can produce a rather large change in publication rates. We show that proposals from high-status authors generate more email conversation among IETF participants. There are several ways that increased early attention might contribute to a greater success in academic publishing, whether directly (getting picked up by a working group) or indirectly (improving the quality of the underlying ideas). Perhaps these "audience building" benefits are an important component of the status effects that sociologists observe in a wide variety of settings. This is an interesting question for future research.

References

Ai, C. & Norton, E. C. (2003). Interaction terms in logit and probit models. *Economics Letters*, *80*(1), 123–129.

Akerlof, G. (1970). Market for lemons - quality uncertainty and market mechanisms. *Quarterly Journal of Economics*, *84*(3), 488–500.

Ayres, I. (2002). Outcome tests of racial disparities in police practices. *Justice Research and Policy*, *4*, 132–142.

Ayres, I. & Waldfogel, J. (1994). A market test for race discrimination in bail setting. *Stanford Law Review*, *46*(5), 987–1047.

Becker, G. S. (1993). Nobel lecture: The economic way of looking at behavior. *The Journal of Political Economy*, *101*(3), 385–409.

Benjamin, B. & Podolny, J. (1999). Status, quality, and social order in the California wine industry. *Administrative Science Quarterly*, *44*(3), 563–589.

Bertrand, M. & Mullainathan, S. (2004). Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *American Economic Review*, *94*(4), 991–1013.

Blank, R. M. (1991). The effects of double-blind versus single-blind reviewing: Experimental evidence from the american economic review. *The American Economic Review*, *81*(5), 1041–1067.

Bothner, M. S., Kim, Y.-K., & Smith, E. B. (2009). How does status affect performance? status as an asset versus status as a liability in the pga and nasc. Working paper.

Burris, V. (2004). The academic caste system: Prestige hierarchies in PhD exchange networks. *American Sociological Review*, *69*(2), 239–264.

DellaVigna, S. (2008). Psychology and economics: Evidence from the field. *Journal of Economic Literature*, *forthcoming*.

Fleming, L. & Waguespack, D. (2009). Scanning the commons? evidence on the benefits to startups of participation in open standards development. *Management Science*, *55*(2), 210–223.

Fuller, B. (2006). Phone interview - november 14 2006.

Gould, R. (2002). The origins of status hierarchies: A formal theory and empirical test. *American Journal of Sociology*, *107*(5), 1143–1178.

Hellerstein, J. K. & Neumark, D. (2004). Production function and wage equation estimation with heterogeneous labor: Evidence from a new matched employer-employee data set. Working Paper 10325, National Bureau of Economic Research.

Knowles, J., Persico, N., & Todd, P. (2001). Racial bias in motor vehicle searches: Theory and evidence. *The Journal of Political Economy*, *109*(1), 203–229.

Manski, C. (2000). Economic analysis of social interactions. *Journal of Economic Perspectives*, *14*, 115–136.

Merton, R. K. (1968). The matthew effect in science. *Science*, *159*, 56–63.

Mouw, T. (2006). Estimating the causal effect of social capital: A review of recent research. *Annual Review of Sociology*, *32*, 79–102.

Peters, D. & Ceci, S. (1982). Peer-review practices of psychological journals: The fate of published articles, submitted again. *The Behavioral and Brain Sciences*, *5*, 187–255.

Phillips, D. & Zuckerman, E. (2001). Middle-status conformity: Theoretical restatement and empirical demonstration in two markets. *American Journal of Sociology*, *107*(2), 379–429.

Podolny, J. (1993). A status based model of market competition. *American Journal of Sociology*, *98*(4), 829–872.

Podolny, J., Stuart, T., & Hannan, M. (1996). Networks, knowledge, and niches: Competition in the worldwide semiconductor industry, 1984-1991. *American Journal of Sociology*, *102*(3), 659–689.

Smart, S. & Waldfogel, J. (1996). A citation-based test for discrimination at economics and finance journals. Working Paper 5460, National Bureau of Economic Research.

Sorenson, O. & Waguespack, D. M. (2006). Social structure and exchange: Self-confirming dynamics in Hollywood. *Administrative Science Quarterly*, *51*(4), 560–589.

Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics*, *87*(3), 355–374.

Stuart, T., Hoang, H., & Hybels, R. (1999). Interorganizational endorsements and the performance of entrepreneurial ventures. *Administrative Science Quarterly*, *44*(2), 315–349.

Washington, M. & Zajac, E. (2005). Status evolution and competition: Theory and evidence. *Academy of Management Journal*, *48*(2), 282–296.

Zuckerman, E. (1999). The categorical imperative: Securities analysts and the illegitimacy discount. *American Journal of Sociology*, *104*(5), 1398–1438.

Tables and Figures

Table 1: Variable Definitions

Variable Name	Definition
Published as RFC	Indicator Variable = 1 if Drafts is published as an RFC
Email Lists	Count of WG listservs where > 1 email mentions the ID
Email Replies	Count of emails on all IETF listservs that mention ID
RFC Cites	Citations from future RFCs to focal RFC
Article Cites	Citations from academic journal articles to focal RFC
Patent Cites	Citations from US patents to focal RFC
Visible Et Al	Indicator: One or more missing authors with "et al" in announcement
Invisible Et Al	Indicator: One or more missing authors without "et al" in announcement
Unlisted Author	Indicator: Visible Et Al or Invisible Et Al
Listed Chair	WG Chair Author is listed in ietf-announce message
Unlisted Chair	WG Chair Author not listed in ietf-announce message
Missing Chair	No WG Chair Author is listed in ietf-announce message
Lead RFCs	Cumulative RFCs published by first authors
Listed RFCs	Cumulative RFCs published visible non-lead authors
Unlisted RFCs	Cumulative RFCs published invisible non-lead authors
Collaboration	Indicator: ID has two or more authors with different affiliations
Missing Collab	Indicator: Collab and announcement only reveals one affiliation
Authors	Number of authors on draft
Intl Authors	Number of non-US authors on draft
Pages	Number of pages in the ID
Draft Queue	Number of IDs submitted in concurrent week
Publication Year	Year when draft is first submitted to IETF

Table 2: Summary Statistics

Variable	Mean	S.D.	Min	Max	N
Published as RFC	0.048	0.21	0.00	1.00	1929
Email Lists	1.389	1.64	0.00	17.00	1929
Email Replies	4.246	8.33	0.00	110.00	1340
RFC Cites	2.602	4.08	0.00	24.00	293
Article Cites	0.548	1.76	0.00	14.00	293
Patent Cites	0.527	1.33	0.00	8.00	293
Visible Et Al	0.262	0.44	0.00	1.00	1929
Invisible Et Al	0.082	0.27	0.00	1.00	1929
Unlisted Author	0.345	0.47	0.00	1.00	1929
Listed Chair	0.100	0.29	0.00	1.00	1929
Unlisted Chair	0.071	0.25	0.00	1.00	1929
Missing Chair	0.068	0.25	0.00	1.00	1929
Lead RFCs	0.564	1.91	0.00	23.00	1929
Listed RFCs	1.789	6.54	0.00	77.00	1929
Unlisted RFCs	0.863	4.59	0.00	76.00	1929
Collaboration	1.150	1.37	0.00	5.00	1929
Missing Collab	0.581	0.49	0.00	1.00	1929
Authors	0.212	0.40	0.00	1.00	1929
Intl Authors	2.930	1.00	2.00	5.00	1929
ID Pages	16.173	16.20	0.00	227.00	1929
Draft Queue	1.590	0.96	0.02	3.79	1929

Table 3: Non-parametric Comparisons

	No WG Chair Authors	Unlisted Non-Lead WG Chair	Listed Non-Lead WG Chair	First Author WG Chair
Individual Drafts				
Drafts (N)	1,606	131	192	287
Published as RFC	0.038	0.046	0.135	0.195
Email Lists	3.906	5.093	6.203	8.851
Email Replies	1.323	1.489	1.870	2.303
Working Group Drafts				
Drafts (N)	383	55	194	242
Published as RFC	0.402	0.455	0.500	0.612
Email Lists	9.238	16.067	15.768	15.982
Email Replies	4.465	4.691	6.649	6.264

Each observation is an Internet Draft submitted to the IETF between 2000 and 2003 with 2 to 5 authors.

Table 4: The Matthew Effect

Unit of Observation = Internet Draft
Dependent Variable = Published as RFC

Model	OLS	Logit†	OLS	OLS	OLS	OLS
Sample	Chair IDs	Chair IDs	All IDs	All IDs	All IDs	Three Authors
Listed Chair	0.132**	0.063**	0.086**	0.087**	0.081**	0.102*
	(0.04)	(0.02)	(0.03)	(0.03)	(0.03)	(0.05)
Unlisted Chair			0.033+	0.036+	0.024	0.043
			(0.02)	(0.02)	(0.02)	(0.05)
Unlisted Author			-0.056**		-0.044**	-0.066**
			(0.01)		(0.01)	(0.02)
Visible Etal				-0.049**		
				(0.01)		
Invisible Etal				-0.076**		
				(0.02)		
Constant	0.008		0.141**	0.157**	0.127*	0.055
	(0.10)		(0.05)	(0.06)	(0.05)	(0.06)
Non-US authors		-0.025*			-0.009**	-0.009
		(0.01)			(0.00)	(0.01)
Multiple Orgs (d)		0.003			0.008	0.018
		(0.02)			(0.01)	(0.02)
Log(Pages)		0.019			0.016**	0.013
		(0.01)			(0.01)	(0.01)
Loq draft queue		-0.090**			-0.051**	-0.030
		(0.03)			(0.01)	(0.03)
Observations	323	323	1,929	1,929	1,929	557
Author Count Effects	Y	Y	Y	Y	Y	Y
Meeting Date Effects	Y	Y	Y	Y	Y	Y
R-squared	0.085		0.040	0.041	0.055	0.102

+10% significance; *5% significance; **1% significance. Robust standard errors.
†Marginal Effects calculated at sample means.

Table 5: What's in a (Missing) Name?

Linear Probability Models of Internet Draft Publication Unit of Observation = Internet Draft Dependent Variable = Published as RFC				
Chair Author	0.085**			0.001
	(0.03)			(0.03)
Unlisted Chair	-0.052+			-0.001
	(0.03)			(0.03)
Log(Lead RFCs)		0.037**		0.037**
		(0.01)		(0.01)
Log(All RFCs)		0.046**		0.046**
		(0.01)		(0.01)
Log(Unlisted RFCs)		-0.025+		-0.025
		(0.01)		(0.02)
Collaboration			0.023+	-0.001
			(0.01)	(0.01)
Invisible Collab			-0.016	-0.004
			(0.02)	(0.02)
Constant	0.160**	0.118**	0.161**	0.119**
	(0.05)	(0.04)	(0.05)	(0.04)
Unlisted Author	-0.048**	-0.031*	-0.048**	-0.029*
	(0.01)	(0.01)	(0.02)	(0.01)
Loq(Draft Queue)	-0.051**	-0.047**	-0.051**	-0.046**
	(0.01)	(0.01)	(0.01)	(0.01)
Obs.	1,929	1,929	1,929	1,929
Author Count Effects	Y	Y	Y	Y
Meeting Date Effects	Y	Y	Y	Y
R-squared	0.05	0.08	0.04	0.08

+10% significance; *5% significance; **1% significance. Robust standard errors in parentheses.

Table 6: Status and Attention

Negative Binomial Regressions

Dependent Variable	Email Lists		Email Replies	
Listed Chair	0.325**	0.175*	0.431**	0.314+
	(0.08)	(0.09)	(0.16)	(0.17)
Unlisted Chair	0.080	-0.097	0.455*	0.341+
	(0.10)	(0.10)	(0.18)	(0.19)
Unlisted Author	-0.258**	-0.231**	-0.478**	-0.483**
	(0.08)	(0.07)	(0.14)	(0.14)
Author(s) with RFCs		0.250**		0.168
		(0.06)		(0.12)
Non-US authors		-0.030		0.006
		(0.02)		(0.04)
Multiple Orgs		0.135**		0.127
		(0.05)		(0.10)
Log(Pages)		0.170**		0.062
		(0.03)		(0.06)
Loq(Draft Queue)		0.157*		0.021
		(0.07)		(0.13)
Constant	0.843**	0.131	1.682**	2.421**
	(0.22)	(0.25)	(0.25)	(0.54)
Observations	1929	1929	1340	1340
Author Count Effects	Y	Y	Y	Y
Year Effects	Y	Y	Y	Y
Log-Likelihood	-3029.35	-2987.38	-3112.95	-3109.74

+10% significance; *5% significance; **1% significance. Robust standard errors in parentheses.

Table 7: Status and Performance

	Negative Binomial Regressions					
Dependent Variable	RFC Cites		Article Cites		Patent Cites	
Chair Author(s)	0.795**	-0.071	0.412	-0.123	0.401	0.601
	(0.19)	(0.23)	(0.41)	(0.52)	(0.43)	(0.60)
Log(RFC Pages+1)	0.558**	0.385**	0.933**	0.836**	0.682*	0.745*
	(0.13)	(0.13)	(0.29)	(0.30)	(0.30)	(0.31)
Multiple Orgs	1.275**	1.006**	1.024+	0.885	-0.043	0.070
	(0.31)	(0.30)	(0.55)	(0.56)	(0.58)	(0.60)
Log(Total RFCs)		0.473**		0.247		-0.153
		(0.09)		(0.20)		(0.22)
Non-US authors		-0.185		-0.236		-0.181
		(0.13)		(0.26)		(0.36)
Log(Versions)		0.231+		0.176		-0.152
		(0.14)		(0.30)		(0.31)
Constant	-0.283	-0.572	-5.580**	-5.756**	-4.624**	-4.402**
	(0.39)	(0.38)	(1.06)	(1.11)	(0.99)	(1.03)
Observations	293	293	293	293	293	293
Author Count Effects	Y	Y	Y	Y	Y	Y
Year Effects	Y	Y	Y	Y	Y	Y
Log-Likelihood	-689.26	-672.68	-186.94	-185.50	-184.17	-183.67

+10% significance; *5% significance; **1% significance. Robust standard errors in parentheses.

Figure 1: The IETF Publication Process

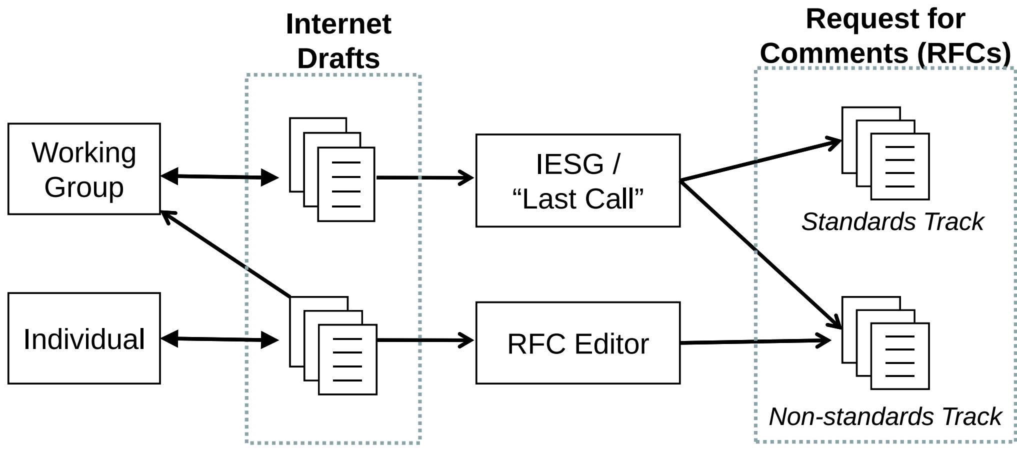

Figure 2: Typical ietf-announce message

An announcement email message (top panel), and the manuscript header and author information section (bottom panel) for an Independent Draft. Thiery Ernst is a Working Group chair.

A New Internet-Draft is available from the on-line Internet-Drafts directories.

Title : Basic Network Mobility Support
Author(s) : R. Wakikawa, K. Uehara, K. Mitsuya, T. Ernst
Filename : draft-wakikawa-nemo-basic-00.txt
Pages : 21
Date : 2003-2-18

This draft proposes a solution for Basic Network Support. It proposes Mobile IPv6 extensions as advocated by the NEMO working group. Our solution differs from Prefix Scope Binding Update …

INTERNET DRAFT
18 Feb 2003

Ryuji Wakikawa
Keisuke Uehara
Koshiro Mitsuya
Thierry Ernst
Keio University and WIDE

Basic Network Mobility Support
draft-wakikawa-nemo-basic-00.txt

… (BODY TEXT) …

Authors' Addresses

Ryuji Wakikawa
Keio University and WIDE
5322 Endo Fujisawa Kanagawa
252-8520 JAPAN
Phone: +81-466-49-1394
Fax: +81-466-49-1395
EMail: ryuji@sfc.wide.ad.jp

Koshiro Mitsuya
Keio University and WIDE
5322 Endo Fujisawa Kanagawa
252-8520 JAPAN
Phone: +81-466-49-1394
Fax: +81-466-49-1395
EMail: mitsuya@sfc.wide.ad.jp

Keisuke Uehara
KEIO University and WIDE
5322 Endo Fujisawa Kanagawa
252-8520 JAPAN
Phone: +81-466-49-1394
Fax: +81-466-49-1395
EMail: kei@wide.ad.jp

Thierry Ernst
Keio University and WIDE
5322 Endo Fujisawa Kanagawa
252-8520 JAPAN
Phone: +81-466-49-1394
Fax: +81-466-49-1395
EMail: ernst@sfc.wide.ad.jp

Figure 3: Et al ietf-announce message

Part B: An announcement email message (top panel), and the manuscript header and author information section (bottom panel) for an Independent Draft. All non-lead author names are suppressed with an "et al" on the email announcement. David Harrington is a working group chair.

A New Internet-Draft is available from the on-line Internet-Drafts directories.

Title : Middlebox Communications (MIDCOM) Protocol Managed Objects
Author(s) : M. Barnes et al.
Filename : draft-barnes-midcom-mib-01.txt
Pages : 16
Date : 2003-7-1

This document describes and defines the managed objects for dynamic configuration of middleboxes. The scope of the middleboxes to which these managed objects apply is limited to NATs and Firewalls …

Internet Draft
Document: draft-barnes-midcom-mib-01.txt

M. Barnes
Nortel Networks
Wes Hardaker
Sparta
D. Harrington
Enterasys Networks
M. Stiemerling
NEC Europe Ltd.
June 2003

Category: Standards Track
Expires: December 2003

Middlebox Communications (MIDCOM) Protocol Managed Objects

… (BODY TEXT) …

Authors' Address

Mary Barnes
Nortel Networks
2380 Performance Drive
Richardson, TX 75082
USA
Phone: 1-972-684-5432
Email: mbarnes@nortelnetworks.com

…..

David Harrington, Co-chair SNMPv3 WG
Enterasys Networks

….

Figure 4: New Internet Draft Submissions (1992-2004)

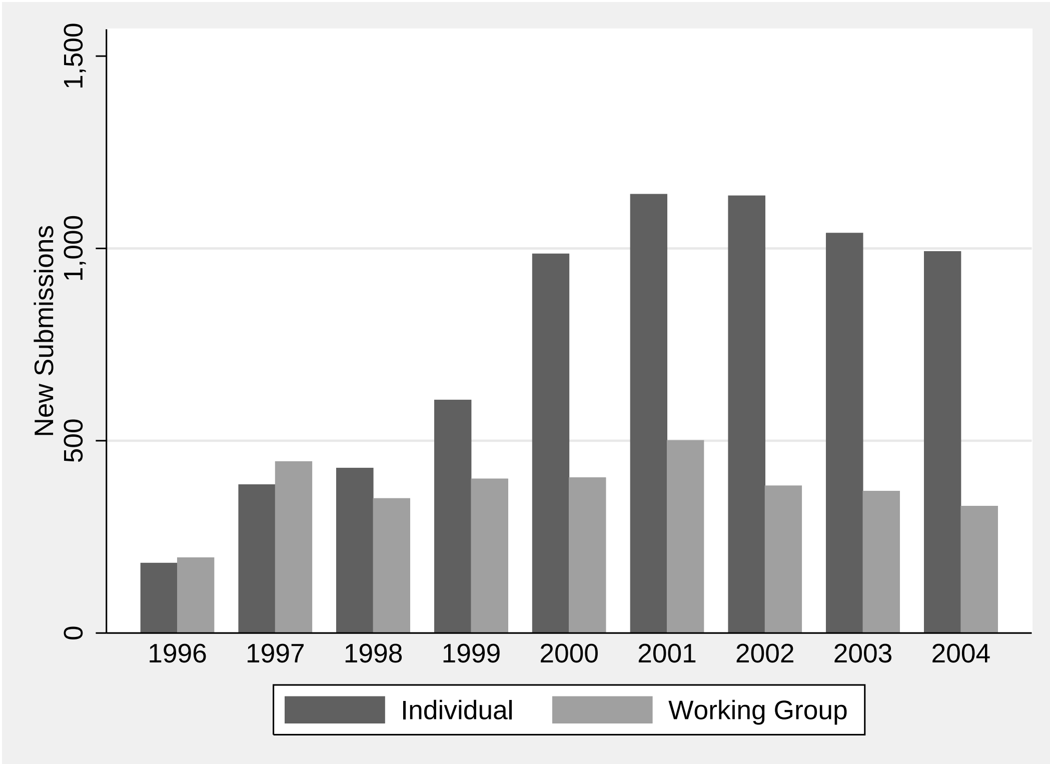

Figure 5: “Et al” Frequency, Meeting Dates, and Draft Submissions