Delayed Returns to Open Source Participation: An Empirical Analysis of the Apache HTTP Server Project*

Il-Horn Hann,
Jeff Roberts,
Sandra Slaughter
Graduate School of Industrial Organization
Carnegie Mellon University

Roy Fielding
Chairman
Apache Software Foundation

First Draft: 03/15/2002

Preliminary and incomplete:
Please do not cite or quote without the authors consent!

Abstract

The availability of commercial quality, free software products such as the Apache HTTP (web) server or the Linux operating system has focused significant attention on the open source development process by which these products were created. One of the more perplexing aspects of open source software projects is why developers freely devote their time and energy to these projects. While many open source participants cite idealistic motives for participation, Lerner and Tirole (2000) argue that developer participation in open source projects may, in part, be explained by existing economic theory regarding career concerns. This research seeks to confirm or disconfirm the existence of economic returns to participation in open source development. Preliminary results of our empirical investigation suggest that greater open source participation per se, as measured in contributions made, does not lead to wage increases. However, a higher status in a merit-based ranking within the Apache Project does lead to significantly higher wages. This suggests that employers do not reward the gain in experience through open source participation as an increase in human capital. The results are also consistent with the notion that a high rank within the Apache Software Foundation is a credible signal of the productive capacity of a programmer.

* We would like to thank Rebecca Tsui for her comments and the numerous open source programmers who have contributed to this study.

Introduction

Open source software development, i.e., public software development projects where participants can read, modify, and redistribute the software source code (OSI 2001), is arguably one of the most exciting phenomena in the software industry today. Open source has played a fundamental role in the development of the Internet by contributing to such remarkable software as TCP/IP, BIND, Sendmail, Linux, and Apache. The open source community has harnessed the Internet like no other by making it the critical piece of its communication and collaboration infrastructure. This prima facie simple innovation has resulted in a revolutionary organization of software production and has sparked discussion on a wide variety of issues, ranging from software development, information architecture, and standards as well as incentives and intellectual property rights.

One widely debated question is why open source programmers contribute voluntarily, thereby foregoing any direct remuneration that they could accrue while working on a commercial system. Often quoted individual level motivations for participating in open source development projects cover a broad spectrum including scratching a “personal itch” with respect to software functionality, enjoyment, and desire to be “part of a team” (Ghosh 1998; O'Reilly 2000; Raymond 2000). Others liken the open source community to a gift culture where the status of a participant depends on ‘what he gives away’ (Raymond 2000) . Alternatively, Lerner and Tirole suggest that open source participation may in part be explained by existing theories of labor economics.

In this paper we investigate whether wages of open source participants are consistent with a signaling incentive. Here, participation in open source projects can work as a signal of the programmer’s productive capacity to current or future employers.

In the following section we describe the argument for delayed incentives in open source participation. We also discuss other factors that might lead to open source participation. Section 3 presents an overview of our research, the Apache HTTP server project, as well as the organization that governs the development process, the Apache Software Foundation. In Section 4, we develop the model used to estimate the delayed returns on open source participation and describe the data set used. The results are presented in Section 5 and we conclude in Section 6.

Explaining Open Source Participation

Motivations for open source participation have been explained from various theoretical perspectives including social psychological, cultural or economic motivations. Eric Raymond, an evangelist of the open source movement, popularized social psychological or cultural explanations of open source participation. In the cultural view, the open source community's truly valuable and protect worthy property is the ownership of ideas or programming projects. Given the abundance of resources, i.e., computing power, bandwidth, and disk space, social status is determined not by what you have, but what you give away. This leads to the 'gift' culture, where the reputation of a programmer is primarily determined by his free contributions (Raymond 2000). As a second explanation, Raymond (2000) offers a 'craftsmanship' model where the artisan aspects of programming motivate developers to create works to be admired not only by themselves but also by others. In both cases developers are motivated through the recognition of their contributions by their peers. Such an explanation finds theoretical support in social psychology (Mauss 1967; Clary, Ridge et al. 1998).

From an economic perspective, a programmer will choose to contribute to an open source project if the benefits outweigh the costs of participation. The primary costs come in the form of opportunity costs for the time spent that could have been otherwise allocated to new or existing projects. Benefits can be categorized as immediate or delayed (Lerner and Tirole 2000). Immediate benefits include the increase of the personal use-value of a product and the satisfaction of having achieved something valuable. Delayed benefits involve the recognition among peers as well as rewards from current or future employers, such as higher wages, stock options or simply more attractive jobs. For both motives, recognition and career concerns, a programmer uses his contributions to signal his capacity to the open source community, to the labor market at large, or even to both.

It is important to point out that some of these different explanations are overlapping. For example, a desire for a higher status within the gift culture may be as strong of an incentive to contribute as career concern incentives. However, as noted by Lerner & Tirole, explaining participation by solely social or cultural factors remains a puzzle for several reasons. First, one could expect to reap similar benefits as part of a commercial software development team obviating the need to participate in an open source project. Second, it is not clear why such noble behavior would be limited to the field of software development (Lerner and Tirole 2000). Moreover, a separation of these motives is, for our purposes, not necessary. As Spence states "A signal is a manipulable attribute or activity which conveys information... in general it is not necessary to insist that the actor, in manipulating the attribute, think of himself as signaling or conveying information" (Spence 1974).

Borrowing further from the labor economics literature, we can distinguish between two approaches to model the value of open source participation: human capital theory and signaling theory. Our data allow us to test both approaches.

Human capital explanations for the value of open source participation are straightforward: Participation allows developers to gain marketable technical skills (Becker 1962; Blaug 1976). This seems an undeniable and obvious benefit of participation. An explanation for open source participation consistent with human capital theory would maintain that open source participation is an investment in training that leads to higher earnings in the future. As an investment, the choice to participate depends upon two considerations. First, the individual considers the opportunity cost associated with participation, and second, the individual considers the expected earnings in the job market after participation. Theory predicts that the greater the investment, the greater the return. Therefore, higher earnings should be correlated with higher levels of open source participation.

While attainment of a skill may be an important result of participation, proponents of a sorting or signaling theory of labor markets argue that participation serves as a signal of individual productive capacities to current and future employers (Weiss 1995). Given a distribution of inherent productivity among potential open source participants, the more productive developers would like to signal their superior productivity to employers (Spence 1973). This is even more important when it comes to software productivity. It is very well known that the productivity difference between an average and a top programmer can be quite large (Weinberg 1998). One study of superstar programmers, for instance, found that the top 1 percent produced 1,272 percent more code than the average. At the same time, due to the nature of programming activities, it might be difficult for a programmer to convey fully his or her productive capacities. While it might be relatively easy to identify the ‘star programmers,’ it is much more difficult to identify above average programmers who have a good understanding of the problem and often develop an efficient solution for the problem at hand. Further, the level of contributions per se might not be the best indicator of productive capacity. Open source projects represent very large-scale, distributed development projects involving thousands of contributions from hundreds of developers (Mockus 2000; O'Reilly 2000). High ability contributors typically make many submissions to the code base, but it is the depth of their understanding, the efficient design of the solution, and their ability to persuade, to get people "on board" with their ideas and strategies that represent the true quality of their contribution. While possible, as a practical matter, it is difficult for employers to efficiently evaluate these qualities based on individual source level contributions. It seems reasonable then that employers seek a reliable proxy that is correlated with these desirable characteristics indicative of or obtained through successful open source participation. If potential employers can use open source participation as a signaling mechanism, then the existence of a “credential” or observable measure of successful participation would allow firms to make inferences about a developer’s productive capacity. In so far as open source participation indicates ability or motivation, it can be used by either employers to screen potential employees or by applicants to signal these desirable traits.

The Apache HTTP Project and the Apache Software Foundation

To determine whether there are economic returns to participation in open source development, we investigate three open source projects under the control of the Apache Software Foundation (ASF). The Apache HTTP (web) server and associated projects are some of the most successful open source products to date. The Apache server, the original ASF project, and its derivatives, have a dominant 63% share of the web server market (Netcraft 2001). Since its inception, the Apache web server has had over 7,000 source code contributions from over 400 different open source developers (Mockus 2000).

The Apache Software Foundation (ASF)

The ASF is a not-for-profit corporation that provides the legal, organizational and financial infrastructure for the software projects gathered under the ASF open-source umbrella. Each of the ASF projects operates autonomously including all aspects of product development. ASF projects are characterized by a "collaborative, consensus-based development process, an open and pragmatic software license, and a desire to create high quality software that leads the way in its field." (Apache 2001). Membership in the ASF is by invitation only and is based on a strict meritocracy. The ASF encompasses seven subprojects related to the development of a full-featured web server product offering. 1) The Apache server project is a freely available source code implementation of an HTTP (Web) server. It is the project around which the Apache Group initially formed. 2) The Apache Portable Runtime project is a free library of C data structures and subroutines designed to facilitate porting the Apache HTTP Server to a host of disparate operating systems. 3) The Jakarta project consists of all Apache related server side Java projects. Jakarta consists of over 18 Java related subprojects. 4) The Apache/Perl project is the integration of the Perl programming language implemented as an Apache HTTP server module. 5) The PHP project is a server side embedded scripting language implemented as an Apache HTTP server module. 6) The Tool Control Language (Tcl) project is an umbrella for Tcl-Apache integration efforts.

These projects combine the Apache web server with the Tcl scripting language. 7) The Apache XML project is home for Apache XML related activities. There are over 9 XML related subprojects.

Although any of the Apache projects could provide an interesting vehicle to explore our research question, we have chosen to concentrate our data collection efforts on the HTTP, Jakarta and Mod_Perl projects for the following two reasons. First, these projects are by far the largest, both in terms of the number of developers and the number of contributions. Second, access to archival data for these projects has proved to be less problematic than for some of the smaller projects.

The Apache "Career"

A common characteristic of open source projects is presence of a strong project leader (Raymond 2000). Apache, however, is unique among open source projects in this regard. Since its inception the Apache project has operated under a model of shared leadership and responsibility (Fielding 1999). This model of shared responsibility is reflected in the principles of the meritocracy that define advancement within the ASF (Apache 2001). As a meritocracy status, responsibility, and benefits are commensurate with contribution. There are five observable levels of recognition or rank within the ASF. In order of increasing status, these are developer, committer, ASF member, project management committee member, and ASF board member. In all cases advancement is in recognition of an individual's commitment and contributions to an Apache project. This hierarchy within the ASF makes the Apache project uniquely positioned to evaluate open source participation. Ideally, data for identifying economic returns to a variable serving as a signal in labor markets would contain exogenous variation in the signal status among individuals with similar levels of human capital (Tyler, Murnane et al. 2000). Participants in ASF projects possess such a variable or credential – their rank or status within the ASF.

Individual reasons for initial involvement in any Apache project vary. Typical reasons cited include reporting a problem or "bug", or fixing a problem in the software that has become a nuisance or impairs usage. Another reason is to extend existing functionality or add new features required by the user or the user’s organization. For the majority of contributors there is a single encounter with the project. Some developers, however, choose a deeper level of involvement and continue to make contributions. If developers’ contributions are significant and consistent over a period of time they may be nominated for an increase in rank from developer to “committer”. The practical significance of attaining the rank of committer on any Apache project is the privilege of submitting code changes directly to the source code repository as opposed to going through an intermediary to have the changes included in the product. An existing ASF member may nominate committers who continue their involvement in the project for ASF membership. ASF membership is largely a matter of recognition and carries with it a certain prestige in the Apache community. ASF members are eligible to be nominated by the ASF Board of Directors or to serve on a project committee. Project committee members are responsible for all aspects of managing an Apache subproject including project plans and roadmaps, release schedules, etc. The ASF Board of Directors makes decisions regarding corporate governance as well as decisions regarding the addition of new projects under the ASF organizational umbrella.

Modeling Delayed Returns to Open Source Participation

Estimating a Wage Equation

We estimate essentially Mincerian wage models that have been traditionally used to test the impact of education on log-earnings (Mincer 1974). The unit of observation is a contributor’s number of contributions, rank and background information in 1999. In our setting, total wage is a lagged function of open source contribution, rank within the Apache Software Foundation, accumulated work experience, accumulated programming skills, education, firm size, firm type (publicly listed or private), job switch, and industry. The general equation estimated is then

$$
\begin{aligned}
\ln\left(TotalWages_{2000}\right) = a &+ b_1 Contribution_{1999} + b_2 Committer^{+}{}_{1999} \\
&+ b_3 \quad WorkExperience_{1999} + b_4 \left(WorkExperience_{1999}\right)^2 \\
&+ b_5 ProgrammingExperience_{1999} + b_6 Masters + b_7 PhD \\
&+ b_8 JobSwitch_{2000} + b_9 FirmSize_{1999} + b_{10} FirmPublic_{1999} \\
&+ b_{12} SWIndustry_{1999} + b_{13} EComIndustry_{1999}
\end{aligned}
$$

The dependent variable is the log of the sum of wages and bonuses. As was noted before, delayed returns of to open source participation may include other benefits, the most important of which are stock options. However, we chose not to ask for detailed information about stock options for the following reasons. First, open source participants working at start-ups told us that they were not allowed to disclose the number of stock options that they were holding, making it impossible to estimate the value of an option. Further, a significant percentage of respondents work at private firms, which makes it difficult to assess the value of these options.

The independent variables are:

- $Contribution_{1999}$: The unit of analysis is commonly known as a "patch". Patches are analogous to modification requests (MRs) in traditional software development environments. Unlike MRs in traditional environments, patches in an open source environment result from largely random developer submissions and have no formal designation or means of tracking. This research follows the method used by Mockus et al. (Mockus 2000) to reconstruct patches from source code archives. For each patch we extract and retain common software metrics including lines of code added and deleted, the date of submission, the names and number of source code files affected by the change, and change log entries. In our equation, $Contribution_{1999}$ measures the number of patches submitted and accepted in 1999. If contributions are a good proxy for the learning experience of an open source programmer, we expect $Contribution_{1999}$ to have a positive effect on wages.

- *Committer*$^{+}_{1999}$: This variable is 1 if the rank of the contributor is 'committer' or above in 1999 and 0 otherwise. The wages of contributors with rank higher than committer were not significantly different from those of committers. If rank is a signal of productive capacity as a programmer, we expect *Committer*$^{+}_{1999}$ to have a positive effect on wages.

- *WorkExperience*$_{1999}$: The total number of years of work experience of a contributor in1999. Typically, we expect that wages increase with work experience, but that the percentage increase declines with higher work experience.

- *ProgrammingExperience*$_{1999}$: The number of years of experience in programming languages and programming tools in 1999. As characteristics such as depth of understanding, design quality of a solution, and programming efficiency rise with experience, this variable proxies for the abilities of the programmer.

- Education: *Masters*/*PhD*: *PhD* is 1 if the contributor has at least 20 years of education and is 0 otherwise. *Masters* is 1 if the contributor has at least17 years and less than 20 years of schooling and is 0 otherwise. Higher education should lead to higher levels of wages.

- *JobSwitch*$_{2000}$: *JobSwitch*$_{2000}$ is 1 if the contributor has switched jobs in 2000 and is 0 otherwise.

- *FirmSize*$_{1999}$: *FirmSize*$_{1999}$ is 1 if the employer of the contributor in 1999 is greater than 200 and is 0 otherwise.

- *FirmPublic*$_{1999}$: *FirmPublic*$_{1999}$ is 1 if the contributor's employer was a public company in 1999 and is 0 otherwise.

- Industry: *SWIndustry*$_{1999}$/*EcomIndustry*$_{1999}$: These variables are set to 1 if the contributor's employer belongs to the software or e-commerce industry respectively, and are set to 0 otherwise.

The Sample and the Data

The data for this research come from two primary sources: Apache project archives and a targeted survey of Apache participants. Archival data are open source project artifacts such as email and source code archives, source code version control meta-data and developer web sites. From these archives, we extracted information pertaining to Apache career advancement as well as individual contributor participation in the development of Apache projects. Survey data came from a questionnaire targeted to Apache contributors. The purpose of this survey was to augment developers' Apache contribution data with their demographic and job history data.

From the archival data, we identified 1,348 contributors. Dr. Roy Fielding, the chairman of the Apache Software Foundation, introduced the survey to the contributors via e-mail in November 2001. 233 e-mails were undeliverable. Of the remaining 1,115 contributors, 325 filled out the survey, yielding a response rate of 29%. For this research, we deleted all contributors who earned income in year 2000 outside of the U.S. or chose not to report any income data. This yielded a sample of 137 contributors.

Results and Interpretations

Table 1 shows descriptive statistics on the ranks and contributions in 1999 of the participants. Of the total of 137 contributors, a large majority, 93, started to contribute in 2000 or later. These people were assigned to the group 'no rank.' The total wages earned in 2000 in the group 'no rank' was statistically not significantly different from the group with rank 'developer.' Our sample contains responses from developers, committers, ASF members, and project committee members, with the max of 172 contributions. Table gives an overview of the variables used in the equation overall and sorted by rank. Overall, open source participants earned total wages of 82K, with a mean of 80K for people with no rank or with rank developers and a mean of 110K for people with rank of committer. People with higher rank earned on average slightly less than committers. In our sample only 39% of all developers contributed in 1999 to open source project, while all committers and 75% of ASF members and project committee members contributed. Our sample is certainly well educated. 11% of the contributors have a doctorate degree and 23% have a master's degree. However, none of the contributors with rank committer and above obtained a Ph.D. and none of the contributors with rank higher than committer have a master degree. Last, Table 3 shows the correlation matrix of the independent variables. As expected, the correlation between rank and contribution is high (0.496) and significant. The correlation between work experience and programming experience is also high (0.524) and significant.

The results of estimating the equation are shown in Table 4. Column 1 shows the effect of contributions, column 2 shows the effect of rank, and column 3 shows the effect of both on the log of wages. The coefficient for $Contribution_{1999}$ is not significant in both specifications (column 1 and column3), indicating that contributions per se do not increase wages. Other specifications such as the sum of contributions until 1999 also did not yield significance. On the other hand, the coefficient for $Committer^{+}_{1999}$ is positive and significant. The wage of contributors with rank committer or above is on average about 29% higher than that of developers after controlling for education, programming skills, work experience, job switch, and firm characteristics (see column 1 and 3). These results suggests that employers of contributors in general do not reward participants for their learning experience in the open source project. However, the higher wage paid to contributors with higher rank is consistent with the idea that the rank conveys sought-after, but typically hard-to-observe characteristics that distinguish above average programmers.

The coefficients of the control variables are generally in the predicted direction. Contributors with a doctorate degree are paid about 20-22% more than college graduates. The coefficient for Master degree holders is, surprisingly, negative, but not significant. Greater work experience increases the salary by 13-14% (significant), but with increasing work experience, the increase in salary is growing more slowly (significant). The coefficient for years of programming experience is not significant, this is due to the high correlation between years of programming experience and years of work experience causing the non-significant result. A job switch increases the salary on average by 15-17% (significant). Public firms also pay more than private firms by about 15-16% more (significant), while firm size is insignificant. Last, contributors in software companies pay significantly more than firms in other industries, about 12-13%.

Conclusion

There are many interesting questions surrounding the open source phenomenon. The research presented here seeks to explore one of the more puzzling aspects regarding open source participation: "Why do developers participate?" We establish two plausible theoretical foundations for the existence of returns to participation from the economics literature, namely, human capital and signaling theory. Our analysis suggests that employers do not reward the accumulation of experience in open source projects. Rather, higher open source rank is associated with higher wages, even when controlling for work experience and programming experience. This is consistent with the notion that firms make inferences about productivity differences based on the rank of the contributor.

Table 1. Descriptive Statistics: Rank and Contribution

Panel A. Ranks

Rank	Frequency	Percent	Cumulative Frequency	Cumulative Percent
No Rank	93	67.88	93	67.88
Developer	36	26.28	129	94.16
Committer	4	2.92	133	97.08
ASF member	1	0.73	134	97.81
Project Committee Member	3	2.19	137	100.00

Panel. B Contributions

$Contribution_{1999}$	Frequency	Percent	Cumulative Frequency	Cumulative Percent
0	116	84.67	116	84.67
1	8	5.84	124	90.51
2	4	2.92	128	93.43
3	2	1.46	130	94.89
5	2	1.46	132	96.35
6	1	0.73	133	97.08
18	1	0.73	134	97.81
20	1	0.73	135	98.54
37	1	0.73	136	99.27
172	1	0.73	137	100.00

Table 2. Descriptive Statistics of the Sample

	Full Sample		Rank = No Rank		Rank = Developer		Rank = Committer		Rank>Committer	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Total Wages	82.15	30.84	80.43	30.19	80.69	31.63	110.00	21.60	107.50	32.02
Salary	78.50	29.14	77.10	28.58	76.53	29.13	107.50	22.17	100.00	34.64
Bonus & Incentives	3.65	9.46	3.33	7.99	4.17	13.17	2.50	5.00	7.50	5.00
$Contribution_{1999}$	2.08	15.15	0.00	0.00	0.78	1.33	49.00	82.44	15.25	16.32
% with $Contribution_{1999}$	0.15	0.36	0.00	0.00	0.39	0.49	1.00	0.00	0.75	0.50
Masters	0.23	0.42	0.27	0.45	0.14	0.35	0.50	0.58	0.00	0.00
Ph.D.	0.11	0.31	0.08	0.27	0.22	0.42	0.00	0.00	0.00	0.00
$ProgrammingExperience_{1999}$	6.67	5.76	6.22	5.77	7.44	5.91	6.50	5.73	10.38	3.30
$WorkExperience_{1999}$	6.28	3.36	5.85	3.55	7.15	2.82	8.29	2.24	6.34	2.65
$JobSwitch_{2000}$	0.16	0.37	0.15	0.36	0.22	0.42	0.00	0.00	0.00	0.00
$FirmSize_{1999}$	0.73	0.45	0.73	0.45	0.78	0.42	0.50	0.58	0.50	0.58
$FirmPublic_{1999}$	0.35	0.48	0.41	0.49	0.17	0.38	0.50	0.58	0.50	0.58
$SWIndustry_{1999}$	0.36	0.48	0.38	0.49	0.28	0.45	0.50	0.58	0.50	0.58
$EcomIndustry_{1999}$	0.23	0.42	0.24	0.43	0.25	0.44	0.00	0.00	0.25	0.50
No. of firm Obs.	137		93 (68%)		36 (26%)		4 (3%)		4 (3%)	

Table 3. Pearson Correlation Matrix[1]

	Contri-bution$_{1999}$	Com-mitter$^{+}_{1999}$	Masters	Ph.D.	Programming Experiene$_{1999}$	WorkEx-perience$_{1999}$	(WorkEx-perience$_{1999}$)^2	JobSwitch$_{2000}$	FirmSize$_{1999}$	Firm-Public$_{1999}$	SW-Industry$_{1999}$	Ecom-Industry$_{1999}$
Contribution$_{1999}$	1.000											
Committer$^{+}_{1999}$	0.496	1.000										
	0.000											
Masters	-0.050	0.010	1.000									
	0.563	0.911										
Ph.D.	-0.037	-0.087	-0.194	1.000								
	0.664	0.310	0.023									
ProgrammingExperiene$_{1999}$	-0.014	0.077	0.020	0.253	1.000							
	0.868	0.372	0.817	0.003								
WorkExperience$_{1999}$	0.088	0.077	-0.047	0.252	0.524	1.000						
	0.304	0.371	0.588	0.003	0.000							
(WorkExperience$_{1999}$)^2	0.080	0.052	0.003	0.276	0.552	0.971	1.000					
	0.353	0.545	0.976	0.001	0.000	0.000						
JobSwitch$_{2000}$	-0.048	-0.109	-0.100	-0.090	-0.023	-0.040	-0.076	1.000				
	0.574	0.205	0.243	0.297	0.788	0.641	0.379					
FirmSize$_{1999}$	0.058	-0.129	0.025	-0.050	-0.042	0.073	0.057	0.087	1.000			
	0.503	0.133	0.772	0.562	0.623	0.400	0.508	0.312				
FirmPublic$_{1999}$	0.141	0.078	0.065	-0.160	0.053	0.011	0.005	0.054	0.447	1.000		
	0.100	0.364	0.453	0.063	0.537	0.903	0.954	0.532	0.000			
SWIndustry$_{1999}$	0.109	0.074	0.056	-0.115	0.071	0.152	0.132	-0.078	-0.095	-0.037	1.000	
	0.205	0.390	0.516	0.180	0.409	0.076	0.123	0.368	0.270	0.665		
EcomIndustry$_{1999}$	-0.073	-0.064	-0.101	0.138	-0.149	-0.092	-0.124	0.087	-0.092	-0.152	-0.412	1.000
	0.399	0.458	0.241	0.108	0.083	0.286	0.150	0.310	0.287	0.076	0.000	

[1] For each pair of correlation, the first row reports the correlation coefficient and the second row the p-value (N=137).

Table 4. Regression Analysis

$$\ln\left(TotalWages_{2000}\right) = a + b_1 Contribution_{1999} + b_2 Committer^{+}{}_{1999} + b_3\ WorkExperience_{1999} + b_4\left(WorkExperience_{1999}\right)^2 + b_5 ProgrammingSkills_{1999} + b_6 Masters + b_7 PhD + b_8 JobSwitch_{2000} + b_9 FirmSize_{1999} + b_{10} FirmPublic_{1999} + b_{12} SWIndustry_{1999} + b_{13} EComIndustry_{1999}$$

	Model 1	Model 2	Model 3
INTERCEPT	10.53*** (97.07)	10.52*** (98.67)	10.52*** (98.22)
$Contribution_{1999}$	0.00 (1.08)	--	-0.00 (-0.09)
$Committer^{+}{}_{1999}$	--	0.29** (2.36)	0.30** (2.08)
Masters	-0.10 (-1.38)	-0.10 (-1.48)	-0.10 (-1.48)
Ph.D.	0.20* (1.95)	0.22** (2.16)	0.22** (2.15)
$ProgrammingExperience_{1999}$	0.00 (0.57)	0.00 (0.38)	0.00 (0.37)
$WorkExperience_{1999}$	0.14*** (3.74)	0.13*** (3.47)	0.13*** (3.45)
$(WorkExperience_{1999})^2$	-0.01** (-2.14)	-0.01* (-1.87)	-0.01* (-1.85)
$JobSwitch_{2000}$	0.15* (1.87)	0.17** (2.12)	0.17** (2.11)
$FirmSize_{1999}$	-0.00 (-0.04)	0.03 (0.37)	0.03 (0.38)
$FirmPublic_{1999}$	0.16** (2.40)	0.15** (2.28)	0.15** (2.27)
$SWIndustry_{1999}$	0.12* (1.83)	0.13* (1.96)	0.13* (1.95)
$EcomIndustry_{1999}$	0.11 (1.44)	0.12 (1.54)	0.12 (1.53)
No. of Obs.	136	136	136
Adjusted R^2	0.40	0.42	0.41

*, **, *** Statistical significance at the 10%, 5%, and 1% level (two-tailed t-test).

References

Apache (2001). http://www.apache.org/foundation/, Apache Software Foundation.

Becker, G. S. (1962). "Investment in Human Capital: A Theoretical Analysis." Journal of Political Economy **70**(5, Part 2: Investment in Human Beings): 9-49.

Blaug, M. (1976). "The Empirical Status of Human Capital Theory: A Slightly Jaundiced Survey." Journal of Economic Literature **14**(3): 827-855.

Clary, E. G., R. D. Ridge, et al. (1998). "Understanding and Assessing the Motivations of Volunteers: A Functional Approach." Journal of Personality and Social Psychology **74**(6): 1516-1530.

Fielding, R. (1999). "Shared Leadership in the Apache Project." Communications of the ACM **42**(4): 2.

Ghosh, R. A. (1998). "Interview with Linus Torvalds: What motivates free software developers?" First Monday **3**(3).

Lerner, J. and J. Tirole (2000). The Simple Economics of Open Source, The National Bureau of Economic Research, Inc. (Accessed: April 4, 2001). http://papers.nber.org/papers/W7600

Mauss, M. (1967). The gift; forms and functions of exchange in archaic societies. New York, Norton.

Mincer, J. (1974). Schooling, experience, and earnings. New York, Columbia University Press.

Mockus (2000). A case study of open source software development: the Apache server. Proceedings of the 22nd international conference on on Software engineering, Limerick Ireland, ACM.

Netcraft (2001). The Netcraft Web Server Survey. Bath, England, Netcraft.

O'Reilly, T. (2000). Open Source: The Model for Collaboration in the Age of the Internet. Computers, Freedom and Privacy, Toronto, Canada.

OSI (2001). The Open Source Definition, The Open Source Initiative. (Accessed: May, 2001). http://opensource.org/docs/definition_plain.html

Raymond, E. (2000). The cathedral and the bazaar, Eric Steven Raymond. (Accessed: October 15, 2000). http://tuxedo.org/~esr/writings/cathedral-bazaar/cathedral-bazaar/

Raymond, E. (2000). Homesteading the Noosphere, Eric Steven Raymond. (Accessed: April 4, 2001). http://tuxedo.org/~esr/writings/homesteading/homesteading/

Spence, M. (1973). "Job Market Signaling." Quarterly Journal of Economics **87**(3): 355-374.

Spence, M. (1974). Market signaling: Information transfer in hiring and related screening processes. Cambridge, MA, Harvard University Press.

Tyler, J. H., R. J. Murnane, et al. (2000). "Estimating the labor market signaling value of the GED." Quarterly Journal of Economics **115**(2): 431-468.

Weinberg, G. (1998). The psychology of computer programming. Silver Anniversary Edition. New York, Van Nostrand Reinhold.

Weiss, A. (1995). "Human-Capital Vs Signaling Explanations of Wages." Journal of Economic Perspectives 9(4): 133-154.