A Large Scale Survey of Motivation in Software Development and Analysis of its Validity

Idan Amit

idan.amit@mail.huji.ac.il

Dror G. Feitelson

feit@cs.huji.ac.il

Department of Computer Science

The Hebrew University, Jerusalem 91904, Israel

Abstract

Context: Motivation is known to improve performance. In software development in particular, there has been considerable interest in the motivation of contributors to open-source.

Objective: We identify 11 motivators from the literature (enjoying programming, ownership of code, learning, self-use, etc.), and evaluate their relative effect on motivation. Since motivation is an internal subjective feeling, we also analyze the validity of the answers.

Method: We conducted a survey with 66 questions on motivation which was completed by 521 developers. Most of the questions used an 11-point scale. We evaluated the answers' validity by comparing related questions, comparing to actual behavior on GitHub, and comparison with the same developer in a follow-up survey.

Results: Validity problems include moderate correlations between answers to related questions, as well as self-promotion and mistakes in the answers. Despite these problems, predictive analysis—investigating how diverse motivators influence the probability of high motivation—provided valuable insights. The correlations between the different motivators are low, implying their independence. High values in all 11 motivators predict increased probability of high motivation. In addition, improvement analysis shows that an increase in most motivators predicts an increase in general motivation.

Conclusions: All 11 motivators indeed support motivation, but only moderately. No single motivator suffices to predict high motivation or motivation improvement, and each motivator sheds light on a different aspect of motivation. Therefore models based on multiple motivators predict *motivation improvement* with up to 94% accuracy, better than any single motivator.

Keywords: Motivation, Software engineering, Open-source development, Survey validity.

1 Introduction

Motivation has a high impact on human performance in many fields [33, 12, 52, 100, 47]. In the context of software development it is especially interesting [20, 83], due to the phenomenon of open-source development [74], where many of the developers are volunteers [82].

We conducted a large-scale survey in order to investigate how developers see the motivators affecting their motivation. Our survey contained questions related to eleven motivators which may affect motivation, taken from prior work. Some of these motivators relate to the culture of open source, such as adhering to its **ideology** [22] and developing software for **self-use** [74]. Other motivators are related to internal reasons, such as **enjoyment** from developing code [39, 11], **learning** from it [16], the feeling of overcoming a **challenge** [63, 59], or gaining **ownership** of a project [52, 15]. Additional motivators extend this to a wider context: being part of a **community** [14, 101], receiving **recognition** for one's work [36, 96] — or, in unpleasant situations, suffering from **hostility** [72, 3]. Finally, there are project-based motivators like a sense of **importance** [40], as well as receiving **payment** for participating in a project [23, 75].

The survey contained 66 questions, covering motivation, satisfaction, demographics, and self-rating of skill. We obtained answers from 1,724 developers, and 521 of them completed the whole survey. A year later we conducted a followup survey, answered by 124 of the original participants. We used the answers to measure the correlation of each motivator with general motivation, evaluate motivators as predictors of general motivation, evaluate how improvement in a motivator improves general motivation, and build a motivation model to measure the overall predictability. Taken together, a motivator that is influential in all these separate ways is likely to have a true impact.

In general, all the motivators were found to contribute to motivation in the predictive sense (knowing of a high motivator means higher probability of motivation). Other than challenge, ideology, and hostility, an improvement in the follow-up answers of the motivator also increases the probability of motivation improvement. On the other hand, none of the motivators is sufficient or necessary on its own for high motivation. Yet, when used together in a predictive model, one can predict well both high motivation and motivation improvement.

Motivation is an internal subjective feeling. When answering regarding our motivation, a valid answer depends on internal identification of the motivation, correct translation into a proper answer, and lack of personal biases [17, 50]. We therefore also analyzed the validity and reliability of the answers. We did this by examining the correlations between questions concerning the same motivator. We also checked the consistency of answers of the same person to the same questions in the original and follow-up surveys. Moreover, we investigated the reliability of the answers, using identification of errors, differences between perception and actual behavior on GitHub, and biases.

We found that the attitudes toward motivators are only moderately stable. The validity also suffers from many kinds of problems. Simple mistakes, like typos, are rather rare. Biases are more common. For example, only 5.6% of

the participants place themselves in a lower-than-neutral level of 'My code is of high quality'. This may indicate that our participants are indeed highly skilled developers. However, the GitHub profiles provided by some participants allowed us to compare their answers to their actual behavior. This showed that the participants overestimated their code quality, their documentation level, and their productivity. Despite these problems, questions belonging to the same motivator are reasonably correlated. Comparing answers of the same developer in the original and follow-up survey also shows a moderate stability. Therefore, the results that were supported by multiple analyses seem to be valid.

This study makes three main contributions:

- We conducted a large-scale survey on software developers' motivation, covering 11 motivators.

- We make several methodological innovations, including about assessing the validity of the results:

  - We asked participants for their GitHub profiles, which enabled comparing survey answers and actual behavior.

  - We conducted a follow-up survey, asking the same people the same questions again after more than a year. This allowed us to measure the answers' stability and the impact of changes in motivators on changes in motivation.

  - We framed the analysis as a supervised learning problem. We initially considered each single motivator as a classifier for high motivation, moved to full models, and then applied the same methods for motivation improvement.

  - We used the correlations of different questions in the same topic to measure absolute and relative coherence.

  - The large scale of the study allowed us to compare answers of different people in the same project, estimating subjectivity.

  - We used several types of analyses in tandem to investigate the relations between motivators and motivation. Most relations were consistent in most or all the methods, testifying to their validity.

- We analyze the results using a machine-learning framework and reach several conclusions regarding motivators and their influence:

  - We corroborate previous work showing that in general motivators from prior work are indeed correlated with motivation.

  - At the same time we find that none of them alone is enough to guarantee motivation, so developers usually need several reasons in order to have high motivation. Also, predictive performance is improved by taking multiple motivators into account.

- Answers regarding motivation have moderate validity, shown by comparing to similar questions, comparing to a different date, or comparing to actual behavior.
- We found that although hostility is rare, when it exists it has a negative influence on motivation. Yet, it tends to be unobserved by others in the same project.

2 Related Work

2.1 Motivation

Motivation, in general and in work context, has been extensively investigated due to its importance. Many theories were suggested. Skinner suggested the operant conditioning, learning behavior due to reward and punishments [86]. Maslow's hierarchy of needs sees self-actualization as the top need [61]. McClelland argues that motivation comes from a mixture of affiliation (society based), authority (opportunities to gain it), and achievements (overcoming challenges) [63]. The equity theory claims that motivation might be hurt due to relative comparison and the feeling of not being fairly treated [3].

In the context of work, the Goal Setting Theory claims that challenging yet achievable goals benefit the motivation [59]. Close in spirit is Vroom's Expectancy Theory [97] that claims that one estimates the outcome, the outcome value, and the probability of the value. Given these, the motivation is determined, and one will have more motivation in tasks where an outcome that is valued is likely to be achieved.

Herzbereg et al. suggested the Motivation-Hygiene Theory [44]. According to it, positive motivation is usually due to intrinsic motivators. However, external hygiene factors might lead to the loss of motivation. Hackman and Oldman suggested the Job Characteristics Theory [40]. They claim that the motivation might come from the job itself, due to the significance, autonomy, skill, identity, and feedback related to the job.

These theories are classical and were introduced long ago. Though they were criticized, they are still beneficial [17].

The relation between job satisfaction and performance is an important research area in organizational psychology [78, 53, 47]. Plenty of empirical work was done in order to understand motivation. Task significance, a motivation cause, was shown to increase productivity [38]. Campbell et al. see performance as a function of motivation but also of knowledge and skill [25].

Job satisfaction also correlated with performance in field studies [100]. Judge et al. performed a meta-analysis and reached 0.3 correlation from 312 studies with 54,417 participants [47].

4

2.2   Motivators

We investigate motivators in software development and specifically in open-source development. Some of these motivators are relevant to any human activity, others to software development or open-source development. Demarco and Lister [29], and also Frangos [33], claim that the important software problems are human and not technological. So, many have investigated motivation in software engineering [55, 20, 34, 37].

Open-source development is the collaborative development of software that is free to use and further modify. The best-known non-technical equivalent is Wikipedia. A seminal description of the phenomenon is given in Raymond's "The Cathedral and the Bazaar" [74]. Payment is probably the most common way to motivate people to perform a task, though it is an extrinsic motivation and therefore its influence is more complex [23, 75]. However, it is common to perform open source software development as a volunteer, which means that salary is not the motivation, making it startling from an economical point of view at first sight [56]. Therefore, the motivation of open source developers was investigated as a specific domain, in an effort to uncover other motivators [101, 96, 28, 57, 42, 76]. In Section 4.2 we list and discuss the 11 motivators used in our study.

2.3   Reliability of Motivation Reports

The limited reliability of motivation reports, a problem that we also cope with, was investigated in prior work. Using self-estimation in a survey might be a threat to the validity of the collected data. There might be biases due to ego defenses[17], the Dunning–Kruger effect [50], subjectivity, and different personal scales. Further, "research on self-esteem (Shavit & Shouval, 1980) [84] has demonstrated empirically that individuals resist lowering favorable self-perceptions" [24]. Previous work has tried to evaluate these difficulties.

Argyle [11] checked the reliability of self-estimation of happiness and showed it is related to peer and supervisory estimation. The Maslach Burnout Inventory validated self-estimation on burnout by comparison with the answers of a close person such as a spouse or a co-worker [60]. Judge et al. [48] also compared a person's and significant other's answers. For work answers "The average correlation between the self and significant-other reports, corrected for unreliability, was r = .68."

Wigert and Harter investigated performance reviews, an area close to motivation [99]. They mention methodological difficulties when one tries to rely on supervisory estimation instead of self-estimation: individual supervisory ratings are a much less reliable measure of performance than objective measures [95], and 62% of the variance in ratings can be attributed to rater bias, while actual performance accounts for just 21% of the variance [81]. Yet, Tsui reports that an employee and his manager's evaluation of effectiveness match [79].

Beatty et al. [18] also compare manager and employee's appraisals. They found that there is agreement on medium performance and some disagreements

on high and low performance. In a second usage there was higher agreement, though it was not clear if it was due to clarification of requirements or just better communication.

As prior work shows, there is a moderate agreement between self-reports and a close person's report. This moderate agreement supports the self-reported answers validity yet warns that they are not perfectly accurate. In this study, we compare the same person's answers to related questions, and the same person's answers in the original and follow-up surveys. Our results are aligned with the prior work, also indicating moderate correlation. We also note that despite all the above concerns, Scott et al. report that Facebook found that surveys are twice more accurate than predictive analytics in employee churn [80].

3   The Survey Instrument

3.1   Design

Our goal is to investigate motivators and motivation. But motivation is not a well-defined concept. To increase validity, we used questions endorsed by prior work and questions that can be compared to actual activity. Some of the motivators were represented by multiple questions, to enable internal coherence validation. This led to the construction of a relatively long survey with 66 questions (see Appendix A). The first section of the survey was "Questions regarding yourself" (18 questions). This part included general questions about motivation and verifiable questions about conduct (e.g., the writing of detailed commit messages). We also asked questions about self-rating of skill.

The second section was "Questions regarding activity in a repository" (28 questions), in which we asked about a specific project and its related behavior. This included our motivation ground truth question: "I regularly have a high level of motivation to contribute to the repository" (based on [65]). It also included questions on eleven motivators from prior work: community, ownership, self-use, etc.

This was followed by the Job Satisfaction Scale questionnaire (10 questions) [45] as is. Job satisfaction is assumed to be related to motivation and its motivators.

In the discussion below we note prior work on each motivator (Section 4.2), and in the replication package we identify the specific source of each survey question [9]. By using questions from prior work, we benefit from a previous validation. Murphy-Hill et al. investigated developer productivity [65]. We took a productivity question from there and also based a motivation question on the same pattern. Amabile et al. developed a survey for intrinsic and extrinsic motivation, from which we took some questions [4]. We took questions regarding hostility from [26] and [49]. Kuusinen at el. investigated flow, intrinsic motivation, and developer experience, from which we took questions on the topic [51].

We also added a "Demography" section (8 questions). Demography is of

6

interest on its own and enables us to compare to prior work. Last, we ask an open question requesting comments Whose goal is to ensure that we did not miss a significant factor in the structured questions.

An important goal of the survey was to enable us to compare answers regarding motivation and actual behavior. For example, we could compare answers to "I write detailed commit messages" and actual commit message length. Therefore, we asked the participants to choose a specific project and provide its name, preferably a public GitHub project. We asked for their GitHub profile for investigating the developer behavior. We also asked for the email from participants who were interested in the research results (and offered them a gift lottery). Email and GitHub profile are personally identifying information, which is usually not collected. We needed them to match the answers to other data related to the same person, but do not include them with the other experimental materials. This was approved by our IRB (study 09032020).

The survey was designed to take about 10 minutes. Most questions used a Likert scale [58, 46] ranging from 1 to 11. All participants saw the sections in the same order, yet the questions order within the sections was randomized. We conducted a pilot to verify the time and make sure that the questions are clear.

3.2 Execution

The survey was conducted using the Qualtrics platform from December 2019 to March 2021, including initial pilots. We obtained 1,724 participants, 521 of them completed the survey.

GitHub is a platform for source control and code development used by millions of users [1]. We initially focused on 1,530 active public GitHub projects with 500+ commits during 2018, described in [7, 8, 5]. About 40,000 developers contributed to these projects that year, of which 9,000 contributed more than 12 commits. We extracted developers' email addresses using the GitHub public email API, fetching the emails of the developers that chose to share them publicly. We sent emails to 3,255 developers with a public email that had enough commits. Due to a mistake, we sent multiple emails to developers that worked on multiple suitable projects, and we apologize for that. We also had a gift card lottery, offering $50 to three of the participants. This channel led to 339 participants, which is 20% of the total.

We also recruited participants in social networks by convenience sampling [30, 2], which led to the remaining 80% of the participants. We used Reddit, an online discussions site, as an important source of participants. Reddit has numerous subreddits, channels dedicated to discussions on specific topics. It has many channels relevant to programmers such as programming language based, operating system based, tools, etc. We posted slowly in different subreddits, to get familiar with the community, as posting at a high rate might be considered as spamming. Each subreddit has different formal and informal rules that should be respected. We found the people showed interest in the survey, which led to discussions and upvotes. These in turn led to more attention to the survey.

As noted above, we asked participants for the name of their project and

their GitHub profile. They provided the names of 484 projects and 303 personal GitHub profiles. But after posting the survey in social media, we noted that many participants stop at the "Questions regarding activity in a **repository**" section since they do not contribute to a GitHub repository. This was also accompanied by direct feedback saying that. Since we were interested in answers about motivation in general and needed the GitHub profile only to link to actual behavior, we changed the questions about "GitHub repository" to "any project", avoiding this drop.

The original survey ended in March 2021. A year after the last response, in April 2022, we sent a follow up survey to the 341 participants that provided their emails in the original survey. We sent them the name of the project on which they initially answered, and asked to answer on the same project in case that they are still active there and on a different project otherwise. The questionnaire was the same as the first one, with the additional validation question "Is it the same project on which you answered last time?". In the follow up survey, 124 out of the 341 participants we reached out to answered (36.3%).

4   Results

We work in the framework of supervised learning, trying to predict a concept using a classifier. The concept, which we try to predict, is high motivation. The classifiers, which provide us with predictions, are the motivators (e.g., high ownership). We evaluate how well motivators predict high motivation, using metrics that compare the prediction to the actual motivation. Interesting metrics are the fraction of those with, say, high ownership who indeed are highly motivated (precision), the improvement over just the prevalence in the population (precision lift), and what fraction of the highly motivated who have high ownership (recall).

The analysis of each individual motivator with respect to general motivation provides simple basic results, yet ignores more complex relations like confounders. We therefore also performed additional, more complex analyses. We analyzed the relations between motivators to see that this risk is small. We used the follow-up survey to analyze motivation improvement of each motivator alone. Last, we built combined models utilising all motivators to avoid risk of confounding and leverage the power of all motivators.

4.1   General Motivation

Productivity is influenced by many factors, including motivation [85]. In our survey, we explicitly asked about the relative influence of motivation and skill. 73% of the participants answered that motivation has more influence on their productivity (answers higher than neutral, see Figure 1), and only 9% answered that their skill is more influential.

Motivation might derive from many motivators, from payment to enjoyment. The concept that we would like to investigate is high motivation to contribute

My motivation has more influence on my productivity, than my skill
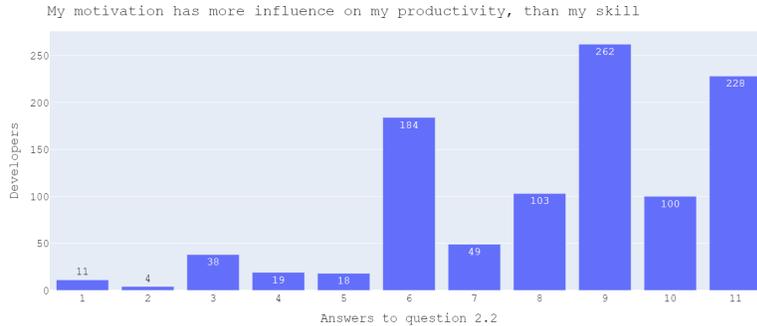
Figure 1: Distribution of answers to motivation vs. skill question.

to a project, and its relations to these various factors.

We measure general motivation using the question "I **regularly** have a **high** level of motivation to contribute **to the repository**" (pattern is based on [65]). The results show that developers are generally motivated (Figure 2). High motivation (at least 9 = 'somewhat agree' on a scale of 1 to 11) was reported by 52.4% of the participants.

I regularly have a high level of motivation to contribute to the repository
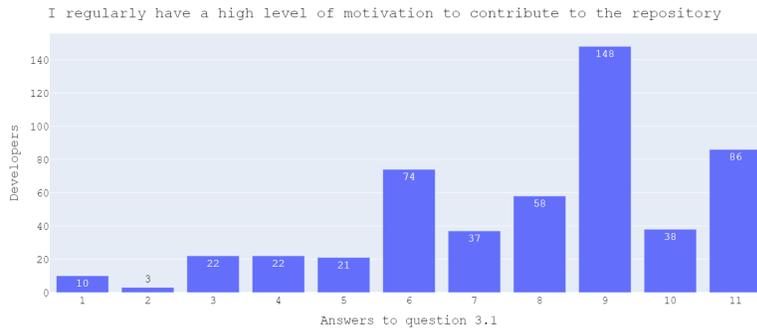
Figure 2: Distribution of answers to general motivation question.

We identified paid people by their answer to a specific question about it (question 3.c). 41% of the participants that answered this question said that they are paid (284 participants). 24% of the participants (377) that had GitHub in their project name were identified as contributing to it (the question encouraged providing a GitHub project when possible). Note that being paid and using GitHub are not mutually exclusive. 38% of participants of GitHub projects that answered the payment question, said that they are paid.

It is generally accepted that motivated workers work longer hours [19]. Showing that our measurement of general motivation exhibits the same relation provides supporting evidence to its validity. And indeed, participants reporting high motivation (at least 9) reported an average of working 19.3 hours a week

on the project, compared to 3.4 hours for those reporting low motivation (below 9). This result may be tainted by mixing data about paid developers with data about volunteers, who are common in open-source projects. We checked this by looking at paid workers separately, out of the participants that reported payment, and the influence of payment is indeed large. For unpaid workers the reported average working hours were 10.8 (high motivation) and 4.5 (low), while for paid workers they were 27.9 (high) and 25.8 (low). Thus, both paid and unpaid participants work more hours when motivated yet the average of unmotivated paid employees developers is higher than that of motivated unpaid developers and therefore we do not mix them.

We planned to also validate the measurement with the similar question from the Job Satisfaction Survey: "Taking everything into consideration, how do you feel about your work?" [45]. However, as discussed in Section 5.5, some of the participants were confused and answered most survey questions on an open source project to which they contribute, yet answered the Job Satisfaction Survey on their regular job. Despite this confusion, the Pearson correlation between the questions is 0.32. When focusing on paid developers, for which the probability of confusion is lower, the correlation is 0.36. While this correlation is moderate, it is similar to the inner correlation in the other motivators as shown below.

4.2   Motivators

The research literature has not produced a canonical agreed list of factors that influence motivation. Mayer et al. reviewed 75 years of motivation measures [62]. This showed that many different factors have an effect, but the agreement between them is limited. We therefore needed to select which ones to include in our study.

We based our list of motivators mainly on Beecham et al.'s review of motivation in software engineering [20] and Gerosa et al.'s [36] work on motivation in open source development. The motivators that we chose have a long history going back to Herzberg's Motivation-Hygiene Theory [44], and therefore were thoroughly investigated over the years (e.g., in general [43] and in software development [27]). Note that we excluded some of the motivators which are less relevant to open-source development, like "Job security" and "Company policies". Conversely, we did include "hostility", which is a *de*motivator (it is common to refer to factors of positive influence as motivators and those of negative influence as demotivators) [3]. Since we have only a single demotivator, we use the term "motivator" to refer to both it and the positive motivators.

Table 1 summarizes the predictive performance of all the motivators. We discuss each of the motivators in the following subsections. Note that in each row we analyzed participants that answered the motivation question and at least one question about the motivator, so populations are not identical.

We use common metrics used in machine learning and information retrieval. The concept that we want to predict is high general motivation. This was operationalized by the answer to the question "I regularly have a high level of

motivation to contribute to the repository" being 9 'somewhat agree' or above.

Usually in machine learning a classification algorithm (e.g., decision tree) is used to create a model, which is a specific rule providing predictions (e.g., if *A & not B*). In contrast, our models are the motivators (e.g., ownership, challenge), also binarized into high and low using 9 as the threshold. Note that in this part the models are pre-defined, and not learnt by a classification algorithm, and we only evaluate their predictive performance.

The cases in which the concept is true are called 'positives' and the positive rate is denoted $P(positive)$ (in our case this is 0.52 as noted above). Cases in which the model is true are called 'hits' and the hit rate is $P(hit)$. For example, a high hit rate for ownership means that many participants report ownership, and we want to see whether they are also generally motivated.

Ideally, hits correspond to positives, but usually some of them differ. Precision, defined as $P(positive|hit)$, measures a model's tendency to avoid false positives (FP). But precision might be high simply since the positive rate is high. Precision lift, defined as $\frac{precision}{P(positive)} - 1 = \frac{P(positive|hit) - P(positive)}{P(positive)}$, copes with this difficulty and measures the *additional* probability of having a true positive relative to the base positive rate. Thus, a useless random model will have precision equal to the positive rate, but a precision lift of 0. Recall, defined as $P(hit|positive)$, measures how many of the positives are also hits; in our case, this is how many of the highly motivated participants also report high ownership.

Table 1:   High Motivation Predictability by Motivator

Motivator	Hit rate	Performance as predictor of motivation			
	(Fraction $\geq$ 9)	Accuracy	Precision	Prec. lift	Recall
Enjoyment	0.74	0.64	0.62	0.18	0.86
Ownership	0.73	0.59	0.57	0.10	0.81
Learning	0.72	0.59	0.58	0.10	0.80
Importance	0.63	0.61	0.61	0.16	0.73
Challenge	0.62	0.63	0.62	0.20	0.74
Self-use	0.56	0.60	0.61	0.17	0.65
Ideology	0.53	0.57	0.59	0.13	0.60
Recognition	0.48	0.58	0.60	0.18	0.56
Payment	0.45	0.55	0.58	0.10	0.49
Community	0.41	0.63	0.67	0.35	0.53
Hostility	0.07	0.52	0.65	0.30	0.08

We now present our analysis of each individual motivator, from the most to the least prevalent. We show how common high answers (9 and above) are to each motivator, in general, for paid developers, and for open-source developers.

4.2.1 Enjoyment

We measure enjoyment [36, 39, 11] by the following questions (numbered by their location in the survey):

2.9 I enjoy software development very much

2.15 I enjoy trying to solve complex problems

3.8 My work on the repository is creative

3.10 I derive satisfaction from working on this repository

An example of the results is shown in Figure 3.

We calculated the average answer to all these questions per participant, and then the average of these averages. This led to an overall average of 9.07. 74% of the participants reported high enjoyment (at least 9 - 'somewhat agree'), more than all other motivators.

76% of the GitHub participants reported high enjoyment and 75% of the paid participants. The correlation of enjoyment with motivation is 0.51, the highest of all motivators.
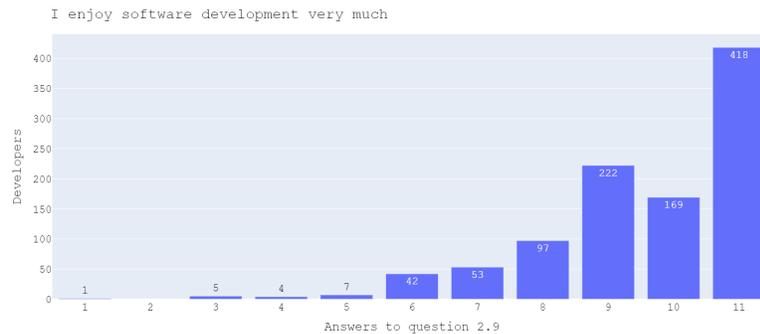
Figure 3: Answers distribution of enjoyment question.

The recall of enjoyment is 0.83, making it very common among people of high motivation, more than both the hit rate and the positive rate. The precision is 0.62 and the precision lift is 0.18, which is moderate.

Note that the hit rate of 74% for enjoyment is significantly higher than the positive rate of 52%. This has a large influence on the predictive metrics. True positives are the intersection of the hits and the positives, so they are bounded by both. But once the hit rate is higher than the positive rate, the precision is bounded from above. In our case, the model hit rate is 74% and the positive rate is 52%, so the model precision can be at most $\frac{52}{74} = 70\%$. The actual precision is 62% which is not high yet is 89% of the bound created from the positive rate and the hit rate.

4.2.2 Ownership

We measure ownership [20, 36, 52, 15, 64] by the following questions:

3.2 - I have complete autonomy in contributing to the repository

3.3 - I have significant influence on the repository

3.4 - I feel responsible to the repository success

3.16 - I am a core member of the repository

The average answer for ownership was 9.02. 73% of the participants reported high ownership (9 or above, second highest of all motivators). This was also the percentage for paid participants; with GitHub participants it was 75%. The correlation of ownership with motivation is 0.24. The recall when predicting high motivation based on high ownership was 0.81, higher than the hit rate. However, the precision is 0.57 and the precision lift is only 0.10, partly since ownership is so common.

4.2.3 Learning

Learning [20, 36, 16] is based on the question:

3.17 - I learn from my contributions

The average of the answers was 9.15, the highest among all motivators. 72% of the participants reported high levels of learning, 70% of the GitHub participants and 77% of the paid ones. The correlation of learning with motivation is 0.23. Learning has a recall of 0.80, indicating that it is another very common characteristic of people with high motivation. Its precision is 0.58 and its precision lift is 0.10, which is relatively low.

4.2.4 Importance

Importance [20, 40, 38] is based on the question:

3.11 - The repository is important

The average was 8.62. The correlation of importance with motivation is 0.35. 63% of the participants report high feeling of importance. The same occurred among GitHub participants, compared to 74% among paid participants. This is rather surprising since we assume that one will have more considerations and constraints in the context of a paid job than in volunteering to open-source projects. So, we would assume that one would have higher freedom to choose by importance when volunteering, leading to a higher rate among GitHub participants, but the data shows the opposite. Importance has precision of 0.61 and precision lift of 0.16. It has a recall of 0.73, which is high in absolute terms and relative to its hit rate.

4.2.5 Challenge

Challenge [20, 63, 59, 89, 71, 83] is based on the question:

3.9 - Working on this repository is challenging

The average of challenge answers was 8.41. 62% of participants reported a high sense of challenge, 60% of the GitHub participants and 66% of paid ones. The correlation of challenge with motivation is 0.30. Challenge has precision of 0.62 and precision lift of 0.20. Its recall is 0.74, which is high in absolute terms and relative to its hit rate.

4.2.6 Self-use

Self-use [36, 74] is based on the question:

3.5 - I'm interested in the repository for my own needs

The average of self-use answers was 7.86. 56% of the participants reported high self-use motivation, 61% of GitHub participants and 43% of paid ones. Note that while usually the probabilities in the entire population, in GitHub, and among paid participants are rather similar, in this case the probabilities are quite different. 'Scratching your own itch' is a well known motivation in open-source [74] so one would expect a higher probability in GitHub. On the other hand, many companies produce organizational software that does not have personal uses, so 43% of paid participants which self-use may sound rather high. The correlation of self-use with motivation is 0.16. Self-use has a recall of almost two thirds, 0.65. This seems to be a unique attribute of open-source, enabling people to develop the software that they need. The precision is 0.61 and the precision lift is 0.17, moderate values. This might be since people see satisfying their need as a task to complete and not an enjoyable activity. Indeed, self-use and enjoyment have a Pearson correlation of just 0.16.

4.2.7 Ideology

Ideology [36, 22] is based on the question:

2.18 - I contribute to open source due to ideology

53% of the participants reported high ideology-based motivation, rising to 61% of GitHub participants, aligned with the ideological roots of open-source development [22]. 49% of paid participants also gave high answers regarding contribution due to ideology. That can be either due to many people being paid to contribute to open-source, or a common habit of paid developers to contribute to open source in their free time. Regardless of the reason, the popularity is surprisingly high. But we note that in the answers to the open question participants said that different ideologies (e.g., 'Software should be free' [88], 'Social good') can lead to contribution to open-source software, and that a finer distinction is needed. The average of ideology answers was 7.34.

The correlation of ideology with motivation is 0.14, the lowest other than for hostility. The precision of ideology is 0.59, the precision list is 0.13, and the recall 0.60.

4.2.8 Recognition

We measure recognition [20, 36, 75, 74, 96] by the following questions:

2.13 I contribute to open source in order to have an online portfolio

2.14 I try to write high quality code because others will see it

3.15 I get recognition due to my contribution to the repository

3.24 In the past year, members of my GitHub community asked questions that show their understanding of my contributions (based on [49])

3.25 In the past year, members of my GitHub community expressed interest in my contributions (based on [49])

The average of all the questions was 7.33, lower than all positive motivators. 48% of the participants reported high recognition-based motivation, 49% of the GitHub ones, and 52% of paid ones. The correlation of recognition with motivation is 0.27. The precision is 0.60 and the precision lift 0.18, indicating a boost to motivation. The recall is 0.56, moderately higher than the recognition hit rate.

4.2.9 Payment

Payment [20, 36, 23, 75] is based on the yes/no question:

3.c - I'm being paid for my work in this repository

We note that remuneration in open-source projects may have many facets. Developers may accrue income from donations or lectures. Their work on the project may help them secure future positions or gain access to future consulting contracts. In the interest of simplicity and precision we use the objective criterion of specifically being paid a salary to define payment.

41% of the participants that answered the payment question said that they are paid. 38% of participants of GitHub projects that answered the payment question said that they are paid.

Payment has precision of 0.58 and precision lift of 0.10, making it one of the weakest motivators in general and the weakest for its hit rate. Payment is also the only motivator that had a negative precision lift when we run the same analysis on the follow-up survey.

Its recall is 0.49, less than the positive rate of high motivation which is 0.52. The correlation of payment with motivation is 0.15, lowest than all but ideology and hostility.

4.2.10   Community

We measure community [20, 36, 63, 21, 14, 101, 32] by the following questions (which we asked to answer only if you are not the only developer in the project):

3.13 Belonging to the community is motivating my work on the repository

3.14 The community is very professional

3.20 The repository's community of developers is more motivated than that of other repositories

3.24 In the past year, members of my GitHub community asked questions that show their understanding of my contributions (based on [49])

3.25 In the past year, members of my GitHub community expressed interest in my contributions (based on [49])

Note the questions 3.24 and 3.25 are about recognition from the community and therefore appear in both motivators.

The average of community answers was 7.36. 40% of the participants reported high community-based motivation, lower than all positive motivators. This is based on 40% of the GitHub participants and 44% of the paid ones. The correlation of community with motivation is 0.42, the second highest. The precision is 0.67 and the precision lift is 0.35, higher than all other motivators. Hence, though the community motivator is not common, when it exists, the probability of high motivation is higher. The recall is 53%, not very high yet 29% higher than the hit rate.

4.2.11   Hostility

Hostility can be viewed as a community with negative influence. Hostility hurts motivation hence it is a demotivator and not a positive motivator. We measure hostility [72, 3, 44, 26] by the following questions (which we asked to answer only if you are not the only developer in the project):

3.6 We have many heated arguments in the community

3.7 I wish that certain developers in the repository will leave

3.22 In the past year, members of my GitHub community put me down or were condescending to me (based on [26])

3.23 In the past year, members of my GitHub community made demeaning or derogatory remarks about me (based on [26], Figure 4)

The average of hostility answers was 2.80. Only 7% of the participants reported high hostility, 4% of the GitHub participants and 7% of the paid ones. The recall is just 8%, yet it is higher than the hit rate. The correlation of

In the past year, members of my GitHub community made demeaning or derogatory remarks about me
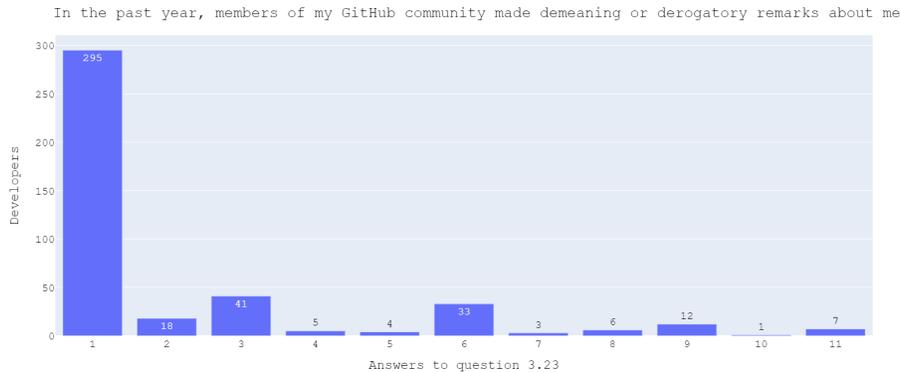
Figure 4:   Answers distribution of a hostility question.

hostility with motivation is 0.01. This is the lowest correlation, close to zero yet not the large negative correlation which is expected.

Surprisingly, hostility also has a relatively high precision of 0.65 and a high precision lift of 0.30. One would expect that knowing that someone suffers from hostility will *reduce* the probability of high motivation, instead of the increase that we see. But the participants who reported high hostility also reported higher averages for all motivators besides payment. A possible explanation is that those participants kept contributing to the project due to the other motivators; those who suffered from hostility and did not have other reasons to stay probably left. Note, however, that we had only 9 people that reported both high hostility and high motivation, so the analysis is not robust.

We identified 10 pairs of developers which contribute to the same project. This allowed us to evaluate their agreement on hostility. Surprisingly, when a person reports heated arguments (question 3.6), the probability that the other participant will agree is just 50%. For the rest of the hostility questions, the other participants never claimed high hostility too. For comparison, in importance and challenge, which also describe aspects of the project, if one participant provided a high half answer (6 or above), the other always agreed. This provides an important indication that hostility might go unnoticed.

4.3   Relations Between Motivators

The above motivators are not all independent. It is therefore interesting to see how they correlate with each other. We now consider motivation and all its motivators as variables and examine the relations between them. We calculated the Pearson correlation between every pair of variables, and looked for connected components on the variables graph, in which an edge exists given a high Pearson correlation.

We use two thresholds of 0.8 and 0.5 Pearson correlations. The aim of the thresholds is to identify strong relations and moderate relations. We compute

the correlations of the variables using three populations: all participants, 284 paid participants (as in Section 4.2.9), and 377 GitHub participants (those who reported a project hosted on GitHub). For each pair of variables, we use the answers of all participants that answered at least one question per variable. The use of the different populations is to see if there are differences in motivation relations in different contexts.
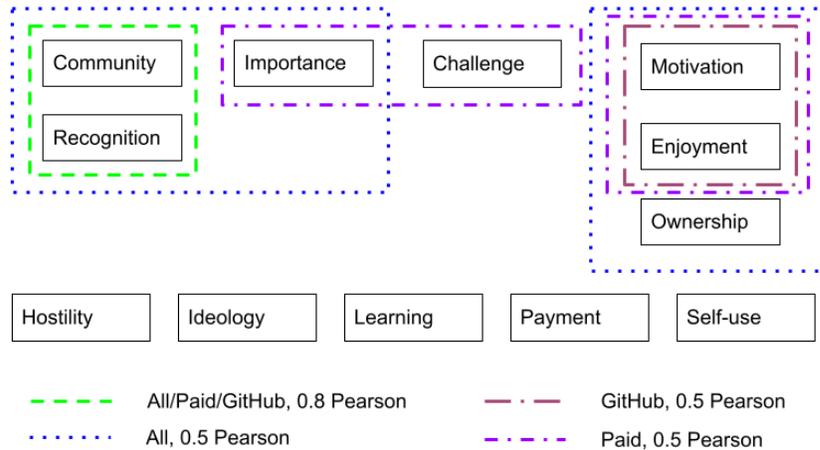
Figure 5: Relations between variables at Pearson levels of 0.8 and 0.5, on the whole population, paid participants, and GitHub participants.

Figure 5 presents the sets of correlated variables identified in different contexts. In all contexts, the 5 variables at the bottom have correlation lower than 0.50. In all 3 populations 'Recognition' and 'Community' are strongly correlated. However, as noted above, recognition and community share two questions regarding recognition from the community (questions 3.24 and 3.25). Removing the common questions, the Pearson correlation of community and recognition is 0.37 on the entire population, 0.35 on GitHub participants, and 0.26 on paid participants. Hence, their actual correlation is much lower.

'Enjoyment' and 'Motivation' are moderately correlated for GitHub users and for paid users, and also correlated with 'Ownership' on all participants. 'Importance' is moderately correlated with 'Recognition' and 'Community' on all participants and with 'Challenge' on paid participants.

In conclusion, the motivator with the highest correlation with motivation is enjoyment, as discussed in Section 4.2.1. Other motivators are less correlated, implying that they are not redundant and each one exposes a different behavioral aspect.

4.4 Motivation Improvement Analysis

An important goal in many data analyses is to uncover causal relations. But causality is hard to define rigorously because it is hard to ascertain that motivator A caused outcome B. The usual approach is to look at correlations between motivators and outcomes in a given dataset, as we did in Section 4.2. We now extend this to look at the *dynamics across time*: the possible correlation between *a change* in a motivator and *a change* in the outcome. For example, we want to see whether an increase in the sense of ownership of a project predicts an increase in motivation.

Such "co-change" analysis [8] is important for the following reason. If causality exists, meaning that in certain contexts motivator A causes outcome B, then a change in A will cause a change in B. But co-change of two motivators does not necessarily imply causality. By identifying instances of co-change, where a change in A correlates with a change in B, we therefore identify cases where causality may be at work.

Note, however, that the actual relationship between motivators and general motivation may be conditioned on other motivators. In this subsection we look at the co-change of motivators and general motivation alone, regardless of context. In the next subsection we consider all the motivators together, to handle cases where the effect of A on B is conditioned on another motivator C.

The change data comes from comparing the original survey and the follow-up survey. In the original survey 341 developers provided their emails. A year after the last response, we reached out and asked them to answer the survey again. We asked them to answer on the same project if they are still active in it. This allowed us to compare the answers of the same person over time. We had 124 follow-up participants in total. 60 of them continued in the same project, and these are the ones we analyze here. For each of them, we look at increases in the motivators and general motivation from one year to the next. Note that if a person reported 3 for ownership in the first survey, and 4 in the follow-up, this is an increase regardless of the values being low.

Table 2: Motivation Improvement Over Time Predictability by Motivator

Motivator	Improvement rate	Prediction of improved motivation			
		Accuracy	Precision	Prec. lift	Recall
Challenge	0.33	0.53	0.10	-0.50	0.17
Ideology	0.30	0.60	0.17	-0.17	0.25
Importance	0.30	0.70	0.33	0.67	0.50
Learning	0.30	0.70	0.33	0.67	0.50
Enjoyment	0.28	0.71	0.34	0.71	0.48
Recognition	0.27	0.72	0.39	0.95	0.47
Self-use	0.22	0.78	0.46	1.31	0.50
Ownership	0.17	0.77	0.45	1.27	0.35
Hostility	0.16	0.70	0.20	-0.01	0.15
Community	0.16	0.75	0.39	0.93	0.28

The probability of improvement in general motivation, i.e. the positive rate, was 20%. Since we included only developers that stayed in the same project for at least a year, probability in the whole population is probably even lower.

The probability of improvement in the different motivators is at most 33% (Table 2). When the precision lift is positive, it tends to be very high. We could not find out why the lift is negative for challenge and ideology. One could expect a larger negative lift for hostility, which did not materialize. This may be explained by developers having other motivators that offset hostility as explained in Section 4.2.11.

The recall is up to 50% for many motivators and higher than their improvement rate (hit rate). This shows that improvement in importance, learning, enjoyment, recognition, and self-use are common when motivation improves.

Only a single person that was not originally paid received a payment in the follow-up, therefore we did not apply co-change analysis to this motivator.

Co-change analysis can be performed in the downward direction too: given a decrease in a motivator, how common is a decrease in general motivation. Results are quite similar to the upward direction and given in the supplementary materials.

The followup survey can also be considered a replication of the original survey. We used all 124 participants that answered the follow-up survey to run the analysis of predicting motivation (as in Table 1). We found a positive precision lift for all motivators besides payment. The agreement supports the results in general. The disagreement in payment indicates that that result is not robust.

4.5   A Combined Motivation Model

All the motivators have a positive precision lift of at least 10% (Table 1). This is aligned with the prior work claiming their positive influence. However, the highest precision lift is just 35%, with 67% precision, for the 'Community' motivator. This means that none of the motivators is a sufficient condition for high motivation or close to it. Hence, a combination of motivators is needed to reach high motivation.

The inputs of machine learning models are named "features". So far, we investigated each motivator as a single feature, ignoring all other motivators. This type of analysis suffers from the threat of confounding variables on one hand, and does not leverage the full power of the data on the other hand. We therefore built combined models to predict motivation, based on all the motivators as features.

The predicted concept was high motivation, operationalized as before by answers of 9 ('somewhat agree') or above to "I regularly have a high level of motivation to contribute to the repository". We had 345 participants that answered this question. The positive rate is 52%.

We used the scikit-learn package for classification algorithms [68]. We used low-capacity small models such as decision trees and logistic regression in order

to obtain simple interpretable models which are also rather robust to overfitting [13]. We also used models of moderate capacity such as random forests, boosting, and neural networks to build models of better representation ability and performance. In order not to bias our results, 70% of the dataset was used for training and performance was evaluated on the 30% remaining test data.

The performance of all the models was rather close, with accuracy ranging from 62% to 77%. Amusingly, the highest accuracy on the test set (77%) was reached by a single node tree, checking high enjoyment. As Table 1 showed, the accuracy when using enjoyment on the whole dataset is just 64%, so this result is accidental. The model with the second highest accuracy (76%) was a neural network [54], whose capacity is high. The simple model of highest accuracy was a logistic regression model [94] which reached accuracy of 72%. Its intercept was -2.04, indicating a general tendency for low motivation. Hostility had a strong negative coefficient of -0.46. All the other motivators had positive coefficients. The highest were enjoyment with 1.13, self-use with 0.61, and importance with 0.59.

Note that models can assign different weights to false positives and false negatives, and trade off precision and recall. Using this, we could build a precision-favoring decision tree model [70] with precision of 81% and recall of 41%. Conversely, a recall-favoring stochastic gradient descent (SGD) model [102] reached recall of 96% with precision of 62%.

We also modeled the co-change dataset of Section 4.4, to predict a change in the motivation based on changes in the motivators as features. Such a model is of interest since assuming that motivation is a function, a co-change model can predict the result of a change.

Two properties of such modeling deserve special attention: accuracy and minimality. We first discuss accuracy. A co-change model allows us to predict *the motivation change* given any *motivator's change*. Perfect accuracy assures us that there are no other external causal variables influencing the samples in our dataset. Assume by contradiction such a variable c, other than the model variables. Hence there is a behavior function g and an assignment of values such that $g(c_1, v_1, ...v_n) \neq g(c_2, v_1, ...v_n)$ where $v_1, ...v_n$ are the values of the model variables. However, since we have perfect prediction given the model variables, it should be that $g(c_1, v_1, ...v_n) = m(v_1, ...v_n) = g(c_2, v_1, ...v_n)$ — a contradiction. Hence such a variable cannot exist. Perfect accuracy is rare, and mostly indicates a problem in the analysis and not capturing all causal variables. However, the accuracy bounds the influence of such external variables.

As for minimality, consider decision trees [70] as models. For each leaf, variables that do not appear in the path to this leaf do not influence the prediction. On the other hand, each variable along the path is necessary, and a change in its value will change the prediction. In this sense, the model is minimal and every variable along the path is required. All the variables that we use here are mutable (can change, in contrast for example to the project creation year). With both perfect accuracy and minimality, each change is explained by the model and the removal of any variable will hurt the prediction of some changes.

When aiming for high accuracy, we built models based on AdaBoost [35] and
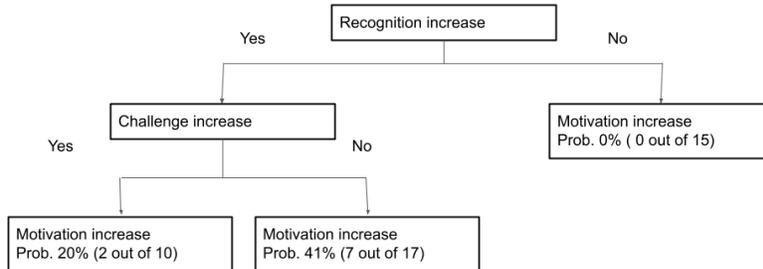
Figure 6:   A decision tree predicting motivation improvement

Neural Networks [54] which reached accuracy of 94%. Aiming for either precision or recall, 100% were reached. The size of these models is high and far from minimality, which is the price for achieving the high accuracy. Alternatively, we found a small model, presented in Figure 6, whose "recognition increase and challenge decrease" leaf reaches 78% recall with 41% precision. Note, however, that the dataset very small dataset with a high VC dimension [93, 92] (due to having many question and wide scale), and therefore the threat of noise is very high. Also, the dataset probably does not fully represent motivation complexity. A larger dataset will probably better represent human motivation behavior but will require a larger model and have lower performance.

5   Analysis of Validity and Reliability

In the previous analysis we analyzed the data as if it is completely reliable. However, the reliability might be limited in many ways. Since the data is given, what we do in this section is to evaluate its reliability from various aspects.

When using the answers of participants, one should check which population they represent. We compare our survey demographics to the demographics of the Stack Overflow survey, answered by around 80 thousand developers worldwide (Section 5.1). We used two channels in order to reach participants: direct emails to developers contributing at GitHub and social media. In Section 5.2 we use machine learning to see how different these populations are. We grouped questions into motivators based on their content, regardless of the answers to them. In Section 5.3 we examine the coherence of the motivators and compare them to grouping based on the answers. The follow-up survey allows us to evaluate the stability of answers, comparing a person's answer in two different dates (Section 5.4). It also provides an additional dataset on which we can check the degree in which our results reproduce. Last but not least, we investigate reliability in the answers themselves, from typos to mistakes and biases (Section 5.5).

5.1 Comparison to the Stack Overflow Survey Demographics

It is hard to define 'developer' (e.g., by education, profession, minimal activity) and therefore to define a representative developers population. However, a possible comparison is provided by the 'Stack Overflow' (SO) annual developer survey. SO is a leading questions and answers website for programming, with more than 100 million monthly visitors. The 2021 survey, overlapping our survey period, was answered by over 80 thousand developers [87]. As far as we know, this is the largest survey of developers and therefore an interesting comparison.

4.2% of our participants identified as females, 95.1% as males, and the rest as others. This is close to the ratio in the 2021 SO survey, which was 5.3%. It is also close to ratios of 4.8% [91] and 11.2% [77] in other sources.

80.7% of the participants in our survey work as professional programmers, compared to 69.7% in SO. The average years of programming experience in our survey (professional or not) is 11.1, representing very experienced developers. Figure 7 shows that in both surveys about half of the participants have at least 10 years of experience.
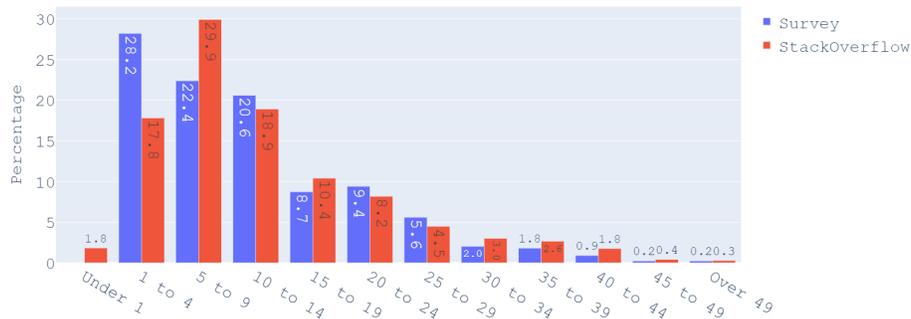
Figure 7: Years of Experience Distribution

Our survey had more participants with high degrees than SO, as shown in Figure 8. The academic domains in our survey were: computer science 46.3%, technology 32.0%, science 9%, business 3.4%, math 2.5%, arts 2.1%, and the rest from other domains.

To conclude, the demographics in our survey are close but not the same as those in Stack Overflow. For example, more people identify as professional developers in our survey. But it seems that the threat of not representing the developers community is low.
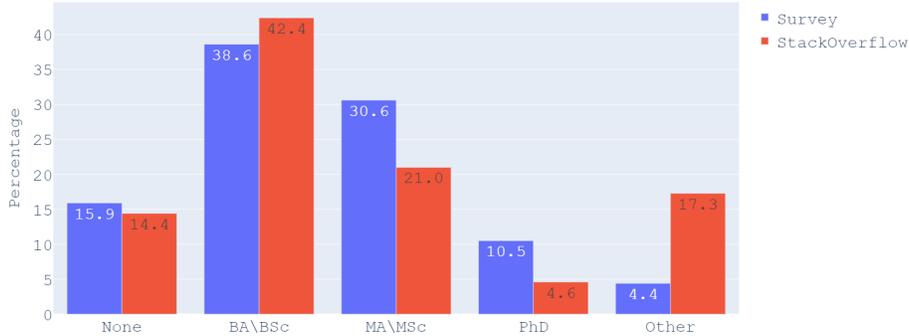
Figure 8: Academic Background

5.2 Differences Between GitHub and Social Media Participants

We received answers from two different populations. 20% of the participants were developers contributing to a GitHub project, which were recruited via direct emails. The other 80% were reached out to in convenience sampling [30, 2], using messages in social media. Differences between the populations might lead to investigating two different behaviors as an averaged one. It is common to compare populations using a comparison of distributions of demographic variables (e.g., age, gender). However, the relevant questions appeared at the end of our survey, and many developers that did not contribute to open source did not reach this part in our first phase. Instead, we reduced the problem of populations difference into a supervised learning problem, trying to predict the source of the participants using the questions in the first part of the survey. The features included open-source specific questions such as 'I contribute to open source in order to become a better programmer' and 'I contribute to open source due to ideology'. Nevertheless, a decision tree model, suitable to a small number of samples, reached an accuracy of only 78%. Even high-capacity models such as SVM or Neural Networks reached an accuracy of only 80%. Note that since the positive rate is 20%, the majority rule betting that all the developers are from social media would also lead to an accuracy of 80%. Such low predictive power does not mean that the populations are similar. However, it means that there is no big obvious difference, even when considering contribution to open source, based on the personal questions. Therefore, we can analyze both populations together, getting a larger dataset and leading to more robust results.

5.3 Internal Coherence of Motivators

The motivators were represented in the survey by one or more questions each. The use of multiple questions (e.g. for 'community') allows us to treat them as labeling functions of the same concept and evaluate their agreement [73, 10]. The agreement, measured by the average Pearson correlation of the related questions, reflects the internal coherence of these motivators. Low coherence might be due to our subjective grouping of questions or due to human nature.

As a reference of the level of correlation that we can expect, we focus on closely related question pairs. For example, 'I am skilled in software development' has a correlation of just 0.62 with 'My code is of high quality'. 'I regularly reach a high level of productivity' and 'I am a relatively productive programmer' have correlation of just 0.57. Table 4 shows that the correlation of the same person answers to the motivation question in the original and follow-up survey is 0.52. Note that a correlation of 0.5 is even higher than the correlation between LOC count and step functions on it [6]. Therefore, coherence of about 0.6 is high.

Table 3:   Motivator Coherence

Motivator	Coherence	Follow-up Coherence
All Questions	0.11	0.07
Community	0.36	0.15
Enjoyment	0.32	0.25
Hostility	0.49	0.60
Ownership	0.58	0.57
Recognition	0.24	0.17

Table 3 presents the coherence of the motivators. 'Coherence' is defined as the average Pearson correlation between all pairs of questions related to the same motivator. The motivators 'Challenge', 'Ideology', 'Importance', 'Learning', 'Payment', and 'Self-use' do not appear in Table 3 since they are based on a single question each, hence our method is not applicable to them. 'Follow-up Coherence' is the same metric as 'Coherence' computed on the follow-up survey. Note that this provides additional support, yet the support is not totally independent since the participants in the follow-up survey also participated in the original one.

The 'All Questions' row represents all the questions together (basically related to motivation), and has rather low coherence. The following motivators all have much higher coherence, indicating that they indeed reflect a meaningful grouping of questions related to specific concepts. The coherence of 'Hostility' and 'Ownership' is relatively high in both surveys, and close to the highest coherence we can expect. The coherence of 'Community', 'Enjoyment', and 'Recognition' is moderate in both surveys.

Next, we compare the grouping of questions as we designed them with an automatic clustering based on their correlations. This will provide additional evidence on whether our grouping was indeed meaningful. The dendrogram
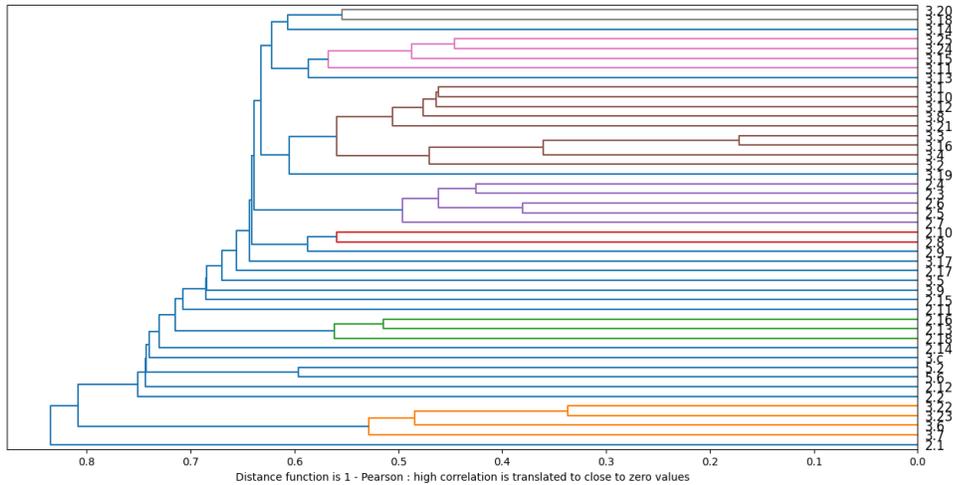
Figure 9: Dendrogram of questions based on their Pearson correlation. See questions text in [9]

in Figure 9 represents a hierarchical clustering [66] of questions based on the correlations between them. The questions 'I have significant influence on the repository' (3.3) and 'I am a core member of the repository' (3.16) are the most correlated questions with Pearson of 0.83 (transformed to $1 - 0.83 = 0.17$ to represent distance in the figure). As we allow weaker correlations, more questions are clustered together, and when we allow Pearson correlation of only 0.3 most questions are already grouped into one big cluster, which is uninteresting. Some of the clusters match our content-based motivators: The orange cluster matches hostility and the lower brown sub-cluster matches ownership. The purple cluster shows a group that we did not consider: productivity and possible productivity improving elements such as good colleagues, physical conditions, and opportunities to use your abilities. Other clusters may overlap our definitions, but also mix in unrelated questions. For example, the pink cluster contains 3 questions about recognition and one about importance, which we feel is not really related. Since our manually built factors are more coherent with respect to content, we use them and not the hierarchical clusters.

5.4 Answers Stability Between Original and Follow-up Surveys

The follow-up survey, conducted one year after the original survey, allowed us to compare the answers of the same person over time. Table 4 shows stability of questions by motivator.

To compare the answers in the two surveys we first compute the Pearson correlation between them. We also compute the differences between them, both the average absolute average difference (column 'Avg. Abs. Diff') and the average relative difference (difference divided by the question value, column 'Avg. Rel. Diff'). 'Pred(25)' [98] is the probability that the follow up answer is in the range of 25% of the initial answer.

Note that the distributions of answers are far from uniform, and some answers are much more popular than others. As a result, there is a high probability for getting the same answer even when the answers are independent. 'Pred(25) Lift' computes the lift, i.e. the extra probability above the expected Pred(25) from two independent answers from the answers distribution.

Pearson correlation, Pred(25), Pred(25) lift, and relative difference indicate stability for almost all motivators. Hostility has a near zero lift and not a large positive one, indicating less stability than expected. Note that the hostility distribution (e.g. Figure 4) has a strong mode in the lowest value, making the independent distribution benchmark very high. Note also that the lift is close to zero hence more likely to be influenced by noise.

Table 4: Similarity of Motivation Type Answers of Same Person in Two Dates

Motivator	Pearson	Avg. Abs. Diff	Avg. Rel. Diff	Pred(25)	Pred(25) Lift
Learning	0.68	0.91	0.04	0.81	0.22
Ownership	0.66	1.02	-0.01	0.83	0.42
Hostility	0.63	1.10	0.38	0.38	-0.02
Enjoyment	0.60	1.06	0.00	0.84	0.29
Ideology	0.57	1.61	0.02	0.74	0.99
Importance	0.54	1.28	0.02	0.74	0.38
Motivation	0.52	1.83	-0.03	0.60	0.30
Challenge	0.51	1.46	0.08	0.70	0.23
Community	0.48	1.46	0.04	0.44	0.04
Recognition	0.45	1.70	0.19	0.54	0.36
Self-use	0.43	2.35	0.04	0.51	0.20

Payment is a binary feature hence its stability should be analyzed with different metrics. The initial and follow-up payment agree in 85% of the cases. 70% of those that were paid in the initial survey were also paid a year later. Only a single person out of the 27 that were not paid in the initial survey got payment in the follow-up.

Note that though the stability is moderate, it is higher than the one reported by Teghtsoonian and Teghtsoonian [90]. There, estimating the stability of the "Stevens's power law exponents" which model human length perception, had correlation of about 0.4 after a few weeks and non-significant after a year. Motivation is more abstract concept than length, more likely to raise ego defenses [17] and yet we showed it is more correlated. Besides, over time the project, the people, and their motivations change [36], which might result in different

answers.

5.5   Face Validity of Answers

To check the validity of the answers in our survey, we looked for mistakes, insincere answers, and biases.

The answer to the gender question was a free text field, in which the participant could write any answer. Only 4 (0.8%) of the answers had a typo (e.g., 'mail', 'boi'). Three of the answers were variants of 'Attack Helicopter', a term "used to disparage transgender people"[1]. Hence, these answers were probably not sincere. In the country question, 1.2% of the answers had a typo.

1.3% of the developers said they had 15 years of experience with GitHub, established in 2008, which was impossible when the survey ended in 2021. A single answer (0.2%) of age of 100 years is probably insincere. Note that these error rates are much better than the 8.5% who seemed to have given a wrong answer to a single simple question in [41], and the 10% failure to identify negatively worded (reverse-coded) items discussed in [69].

The job satisfaction questions were taken from a survey of 9,900 Australian clinical medical workers published in 2011 [45]. Amusingly, software developers were on average less satisfied in all questions. More importantly, questions about payment are irrelevant to volunteers and questions about community are irrelevant to people working alone. We explicitly asked to skip these questions if they are irrelevant. However, 57.7% of the people that answered that they are not paid, answered the payment satisfaction question. Therefore, it seems they answered regarding their salary from a different job, not related to the discussed project. Some participants made comments in the open question that support this.

Some developers answered that they work on a public GitHub project. For these projects, we checked the number of developers who committed code. Of five developers who were found to work on single person projects, one answered most of the community-related questions, which are irrelevant to such projects.

There were questions in which the change in the follow-up survey is known in advance: age and experience should grow linearly with time. The follow-up question was about a year after the first one. In order to avoid rounding mistakes (e.g., a 20.5 years old participant might answer either 20 or 21), we consider answers as "unreasonable" only if the follow up answer was more than a year lower, or at least three years higher. 26% of the answers about experience exhibited such unreasonable differences. Two participants lost 5 years of experience each, somehow compensated by a participant that gained 11 years of experience in about a year. 16% of the answers regarding GitHub experience were unreasonable. However, for age, which has a higher presence in daily life, there were no unreasonable differences.

It seems that the biggest reliability problem comes from human failings [69], bias due to ego defenses[17], or the Dunning–Kruger effect (that people with

---

[1]https://en.wikipedia.org/wiki/I_Sexually_Identify_as_an_Attack_Helicopter

lower capabilities tend to have higher self-esteem) [50]. Only 5.6% of the participants gave a low answer to 'My code is of high quality', going up to 20.8% when including neutral answers (6 on the 11-point scale). The Pearson correlation with years of experience, a common method to estimate skill [31], was a very low 0.06. Moreover, first degree holders gave answers averaging 9.05, higher than all others. People trained in computer science gave answers averaging 9.3, lower than the 10.5 average in math, yet a bit higher than arts (9.0), science (8.7), technology (8.6), and business (7.5).

Using the participants' GitHub profile, we can compare their actual activity to their self-perception. People that answered that they write detailed commit messages (at least 9 - 'somewhat agree'), had average commit message length of 89 characters, placing them in the 61 percentile of GitHub developers, not very far from the median. Participants saying that they write high quality code have corrective commit probability (CCP) [8] of 0.36 (investing more than a third of their work in bug fixing), worse than 81% of the GitHub developers. We measured productivity by commits per working day [65, 67]. Developers that consider themselves to be productive contribute 3.14 commits per working day on average. This is lower than the 3.34 commits of developers that do not consider themselves to be relatively productive.

It seems that there is also a bias leading to higher answers about the participant than about the community. The average answer for questions about themselves is 9.1, 24% more than the average answer to questions about the project. A somewhat smaller difference of 4.5% to 17.1% was found when the questions were essentially paired (e.g, 'My code is of high quality' and 'The quality of the code in this repository is better than others').

We cannot accurately aggregate the probability of mistakes. The probability of identifying an insincere answer is low, around 0.2%. Mistakes typically occur in few percent of the answers, yet in specific questions (the job satisfaction in our case) might be around 50%. Biases are the largest threat to validity, as demonstrated by the 79% participants that consider their code to be of high quality. Only 5.6% of the participants gave a low answer to 'My code is of high quality', and 20.8% when including neutral. The 5.6% that think that their code is of low quality seem to be either more modest or more realistic.

6   Threats to Validity

Motivation, and motivators, are not well defined. Therefore, it is hard to measure them or even evaluate how well a measurement method performs. We cope with this threat using several methods like using questions from prior work, which were already considered to be useful.

The selection of motivators and questions has subjective aspects, and others could be chosen. We based our selection on motivators with massive prior work in motivation in general, in software development, and open source.

Some questions have systematic problems. The job satisfaction questions were answered by many participants on their day job. In self-assessment ques-

tions developers have a very high perception of themselves, not aligned with their actual performance (Section 5.5). We therefore avoided using these answers for motivation analysis.

In order to further reduce the influence of individual questions, we grouped questions by motivators. While we still do the analysis at the question level too (available in the supplementary materials), the aggregation reduces the weight of a specific answer and makes the concepts more robust. However, answers to different questions on the same concepts are only moderately correlated (Section 4.2), so one can argue that our grouping is not correct. Indeed, we grouped questions by subjective judgment of their contents, and in principle a different taxonomy could be used. We compared our content based grouping to the one inferred from correlation (Section 4.2). The match is only partial hence our grouping is supported yet there are also different justifiable groupings.

Surveys are answered by people. Answers of the same person change over time (Section 5.4) and therefore analysis based on the original survey might not agree with the same analysis on the follow-up survey. Concerning the concept of hostility, different people provided widely different answers about the same project (Section 4.2.11). In this case this does not represent a data-quality problem, since the answers actually represent different experiences, and the difference itself is an important result.

We measured the relations between motivators and motivation in multiple ways: correlation, predictive performance, and co-change. A similar result in all methods (e.g., community increases motivation by 20%) would have been very reliable. However, there are many quantitative and even several qualitative differences in the results. For example, Table 1 shows that all motivators have positive precision lift in high motivation prediction, hence knowing of a positive motivator increases the probability of high motivation. On the other hand, in the follow up analysis presented in Table 2, three of the eleven motivators have negative precision lift, hence knowing of their increase from the original survey indicates higher probability of motivation reduction. While negative lift is expected for hostility, the results for challenge and ideology disagree with the high motivation prediction.

Throughout this research we obtained many results. While our number of participants is very high for a survey, we analyzed the answers in many ways. In some scenarios (e.g., developers in the same project, developers answering the follow up), the numbers are quite small. Statistical learning theory [93, 92] tells us that in such cases several of the empirical results will probably be different from the actual ones. This is an inherited threat from the dataset size and analysis type, which should be resolved by replication studies obtaining more data and supporting the results in different analyses.

Similarity between the participants group and the desired population increases the probability of generalization. We discussed in the demographics section (Section 5.1) that though our participants resemble the Stack Overflow survey participants (while being somewhat more professional), it is not clear what is the general developers group.

7   Conclusions

We conducted a large survey of software developers regarding motivation. We grouped the questions by motivators and analyzed their relation to motivation.

Our supervised-learning-based analysis of motivators ranged from using each motivator as a classifier for motivation, through using motivator improvement as a classifier for motivation improvement, to constructing models that use the combined power of multiple motivators for improved prediction. We confirmed that previously suggested motivators do indeed contribute to motivation. At the same time, the influence of each individual motivator is limited, as also noted in prior work [43, 27].

Apparently, the motivation of different developers, working in different contexts, may be influenced by different motivators. No single motivator by itself is sufficient for inducing high motivation. At the same time, none of the motivators is strictly necessary. An analysis of the relations between them indicated that motivators tend to have low correlation. This indicates that one should not look at motivators from the prism of which is the "most important" one; a better description is that each one of them captures a different aspect of motivation [62], and multiple aspects should be satisfied in order to have high motivation.

All motivators have coherence higher than the set of all questions together, but only hostility and ownership have rather high coherence. In general, all motivators are at least moderately coherent and predictive, in all analyses. However, out of the eleven motivators, ten motivators (excluding ownership) did not meet all three criteria of high coherence, stability, and predictive power. This indicates that the bar that we set is high.

It is also interesting to notice the relative position of payment. Trying to predict high motivation based on a single motivator, payment has precision lift of 10%, the lowest value of all positive motivators, and the only negative lift on the follow-up survey. Since payment is the common way to promote motivation in businesses, it is important to note that other motivators might lead to a larger effect.

Hostility is a very coherent demotivator. However, different people in the same project disagree on hostility, implying that it is not noticed by others. Hence, not noticing hostility is not enough to assure lack of hostility in a project, and therefore actively looking for it might be needed.

Participants who reported an improvement in the interest expressed in them had a large tendency for improvement in motivation. Recognition, and specifically expressing interest, is free, applicable in all situations, and influential. Considerate behavior and looking for practical benefits coincide here. Be kind and give recognition, it is likely to pay off.

Data Availability

All experimental materials (except for identifying data such as emails and GitHub profiles) is available at [9].

Acknowledgments

References

[1] Github user search showing 37,446,292 available users. URL https://github.com/search?q=type:user&type=Users. Accessed: 2019-05-23.

[2] A. S. Acharya, A. Prakash, P. Saxena, and A. Nigam. Sampling: Why and how of it. *Indian Journal of Medical Specialties*, 4(2):330–333, 2013. doi:10.7713/ijms.2013.0032.

[3] J. S. Adams. Toward an understanding of inequity. *Journal of Abnormal Psychology*, 67(5):422–436, 1963. doi:10.1037/h0040968.

[4] T. M. Amabile, K. G. Hill, B. A. Hennessey, and E. M. Tighe. The work preference inventory: Assessing intrinsic and extrinsic motivational orientations. *Journal of Personality and Social Psychology*, 66(5):950–967, 1994. doi:10.1037/0022-3514.66.5.950.

[5] I. Amit. End to end software engineering research. 2021, arXiv:2112.11858 [cs.SE].

[6] I. Amit, N. B. Ezra, and D. G. Feitelson. Follow your nose – which code smells are worth chasing? 2021, arXiv:2103.01861 [cs.SE].

[7] I. Amit and D. G. Feitelson. Which refactoring reduces bug rate? In *Proceedings of the 15th International Conference on Predictive Models and Data Analytics in Software Engineering*, PROMISE'19, pages 12–15, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3345629.3345631.

[8] I. Amit and D. G. Feitelson. Corrective commit probability: A measure of the effort invested in bug fixing. *Software Quality Journal*, 29(4):817–861, Aug 2021. doi:10.1007/s11219-021-09564-z.

[9] I. Amit and D. G. Feitelson. Replication package https://github.com/evidencebp/motivation-survey, Dec 2023.

[10] I. Amit, E. Firstenberg, and Y. Meshi. Framework for semi-supervised learning when no labeled data is given. U.S. patent #US11468358B2, 2017. URL https://patents.google.com/patent/US11468358B2.

[11] M. Argyle. Do happy workers work harder? the effect of job satisfaction on job performance. In R. Veenhoven, editor, *How harmful is happiness? Consequences of enjoying life or not*, pages 94–105. Universitaire Pers, Rotterdam, The Netherlands, 1989.

[12] D. Ariely, E. Kamenica, and D. Prelec. Man's search for meaning: The case of legos. *Journal of Economic Behavior & Organization*, 67(3-4):671–677, 2008. doi:10.1016/j.jebo.2008.01.004.

[13] D. Arpit, S. Jastrzębski, N. Ballas, D. Krueger, E. Bengio, M. S. Kanwal, T. Maharaj, A. Fischer, A. Courville, Y. Bengio, and S. Lacoste-Julien. A closer look at memorization in deep networks. In *34th International Conference on Machine Learning*, pages 233–242. PMLR, 2017.

[14] G. Asproni. Motivation, teamwork, and agile development. *Agile Times*, 4, Mar 2004.

[15] N. Baddoo and T. Hall. Motivators of software process improvement: an analysis of practitioners' views. *Journal of Systems and Software*, 62(2):85–96, 2002. doi:10.1016/S0164-1212(01)00125-X.

[16] K. M. Bartol and D. C. Martin. Managing information systems personnel: A review of the literature and managerial implications. *MIS Quarterly*, 6(4):49–70, Dec 1982. URL http://dl.acm.org/citation.cfm?id=2025521.2025525.

[17] N. Bassett-Jones and G. C. Lloyd. Does herzberg's motivation theory have staying power? *Journal of Management Development*, 24:929–943, Dec 2005. doi:10.1108/02621710510627064.

[18] R. W. Beatty, C. E. Schneier, and J. R. Beatty. An empirical investigation of perceptions of ratee behavior frequency and ratee behavior change using behavioral expectation scales (BES). *Personnel Psychology*, 30(4):647–658, 1977. doi:10.1111/j.1744-6570.1977.tb02333.x.

[19] D. G. Beckers, D. van der Linden, P. G. Smulders, M. A. Kompier, M. J. van Veldhoven, and N. W. van Yperen. Working overtime hours: relations with fatigue, work motivation, and the quality of work. *Journal of Occupational and Environmental Medicine*, pages 1282–1289, 2004.

[20] S. Beecham, N. Baddoo, T. Hall, H. Robinson, and H. Sharp. Motivation in software engineering: A systematic literature review. *Information & Software Technology*, 50(9-10):860–878, 2008. doi:10.1016/j.infsof.2007.09.004.

[21] J. K. Beggan. On the social nature of nonsocial perception: The mere ownership effect. *Journal of Personality and Social Psychology*, 62(2):229, 1992. doi:10.1037/0022-3514.62.2.229.

[22] S. Belenzon and M. Schankerman. Motivation and Sorting in Open Source Software Innovation. CEP Discussion Papers dp0893, Centre for Economic Performance, LSE, Oct 2008. URL https://ideas.repec.org/p/cep/cepdps/dp0893.html.

[23] R. Bénabou and J. Tirole. Intrinsic and extrinsic motivation. *The Review of Economic Studies*, 70(3):489–520, 2003.

[24] D. J. Campbell and C. Lee. Self-appraisal in performance evaluation: Development versus evaluation. *The Academy of Management Review*, 13(2):302–314, 1988. doi:10.2307/258579.

[25] J. P. Campbell, R. A. McCloy, S. H. Oppler, and C. E. Sager. A theory of performance. In N. Schmitt, W. C. Borman, and Associates, editors, *Personnel Selection in Organizations*, pages 35–70. Jossey-Bass Pub., 1993.

[26] L. Cortina, V. Magley, J. Williams, and R. Langhout. Incivility in the workplace: Incidence and impact. *Journal of Occupational Health Psychology*, 6:64–80, Feb 2001. doi:10.1037/1076-8998.6.1.64.

[27] J. D. Couger. Motivators vs. demotivators in the is environment. *Journal of Systems Management*, 39(6):36, 1988.

[28] A. B. Cristina and C. Rossi. Altruistic individuals, selfish firms? the structure of motivation in open source software. *First Monday*, 9:9, 2004. doi:10.5210/fm.v9i1.1113.

[29] T. DeMarco and T. Lister. *Peopleware: Productive Projects and Teams*. Dorset House, 1987.

[30] I. Etikan, S. A. Musa, R. S. Alkassim, et al. Comparison of convenience sampling and purposive sampling. *American Journal of Theoretical and Applied Statistics*, 5(1):1–4, 2016. doi:10.11648/j.ajtas.20160501.11.

[31] J. Feigenspan, C. Kästner, J. Liebig, S. Apel, and S. Hanenberg. Measuring programming experience. In *20th IEEE International Conference on Program Comprehension*, pages 73–82, 2012. doi:10.1109/ICPC.2012.6240511.

[32] F. Ferreira, L. L. Silva, and M. T. Valente. Turnover in open-source projects: The case of core developers. In *Proceedings of the 34th Brazilian Symposium on Software Engineering*, pages 447–456, 2020. doi:10.1145/3422392.3422433.

[33] S. A. Frangos. Motivated humans for reliable software products. In D. Gritzalis, editor, *Reliability, Quality and Safety of Software-Intensive Systems: IFIP TC5 WG5.4 3rd International Conference on Reliability, Quality and Safety of Software-Intensive Systems*, pages 83–91. Springer US, Boston, MA, May 1997. doi:10.1007/978-0-387-35097-4_7.

[34] C. França, F. Q. B. da Silva, and H. Sharp. Motivation and satisfaction of software engineers. *IEEE Transactions on Software Engineering*, 46(2):118–140, Feb 2020. doi:10.1109/TSE.2018.2842201.

[35] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997. doi:10.1006/jcss.1997.1504.

[36] M. Gerosa, I. Wiese, B. Trinkenreich, G. Link, G. Robles, C. Treude, I. Steinmacher, and A. Sarma. The shifting sands of motivation: Revisiting what drives contributors in open source. In *Proceedings of the 43rd International Conference on Software Engineering*, pages 1046–1058, 2021. doi:10.1109/ICSE43902.2021.00098.

[37] S. A. K. Ghayyur, S. Ahmed, S. Ullah, and W. Ahmed. The impact of motivator and demotivator factors on agile software development. *International Journal of Advanced Computer Science and Applications*, 9(7), 2018. doi:10.14569/IJACSA.2018.090712.

[38] A. Grant. The significance of task significance: Job performance effects, relational mechanisms, and boundary conditions. *The Journal of Applied Psychology*, 93:108–24, Feb 2008. doi:10.1037/0021-9010.93.1.108.

[39] D. Graziotin, F. Fagerholm, X. Wang, and P. Abrahamsson. What happens when software developers are (un)happy. *Journal of Systems and Software*, 140:32–47, 2018. doi:10.1016/j.jss.2018.02.041.

[40] J. R. Hackman and G. R. Oldman. *Motivation Through the Design of Work: Test of a Theory*. Academic Press, New York, 1976.

[41] S. Herbold, A. Trautsch, B. Ledel, A. Aghamohammadi, T. A. Ghaleb, K. K. Chahal, T. Bossenmaier, B. Nagaria, P. Makedonski, M. N. Ahmadabadi, et al. A fine-grained data set and analysis of tangling in bug fixing commits. *Empirical Software Engineering*, 27(6):1–49, 2022. doi:10.1007/s10664-021-10083-5.

[42] G. Hertel, S. Niedner, and S. Herrmann. Motivation of software developers in open source projects: an internet-based survey of contributors to the linux kernel. *Research Policy*, 32(7):1159–1177, 2003. doi:10.1016/S0048-7333(03)00047-7. Open Source Software Development.

[43] F. Herzberg. One more time: How do you motivate employees? *New York: The Leader Manager*, pages 433–448, 1986.

[44] F. Herzberg, B. Mausner, and B. B. Snyderman. *Motivation to Work*. Wiley, New York, 1959.

[45] D. Hills, C. Joyce, and J. Humphreys. Validation of a job satisfaction scale in the australian clinical medical workforce. *Evaluation & the Health Professions*, 35:47–76, Mar 2011. doi:10.1177/0163278710397339.

[46] A. Joshi, S. Kale, S. Chandel, and D. K. Pal. Likert scale: Explored and explained. *British Journal of Applied Science & Technology*, 7(4):396, 2015. doi:10.9734/BJAST/2015/14975.

[47] T. Judge, C. Thoresen, J. Bono, and G. Patton. The job satisfaction-job performance relationship: a qualitative and quantitative review. *Psychological Bulletin*, 127:376–407, Jan 2001. doi:10.1037/0033-2909.127.3.376.

[48] T. A. Judge, E. A. Locke, C. C. Durham, and A. N. Kluger. Dispositional effects on job and life satisfaction: The role of core evaluations. *Journal of Applied Psychology*, 83:17–34, 1998. doi:10.1037/0021-9010.83.1.17.

[49] A. Kluger and O. Bouskila-Yam. Facilitating listening scale (FLS). In D. L. Worthington and G. D. Bodie, editors, *The Sourcebook of Listening Research: Methodology and Measures*, pages 272–280. Wiley Online Library, Aug 2017. doi:10.1002/9781119102991.ch25.

[50] J. Kruger and D. Dunning. Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology*, 77:1121–1134, 1999. doi:10.1037/0022-3514.77.6.1121.

[51] K. Kuusinen, H. Petrie, F. Fagerholm, and T. Mikkonen. Flow, intrinsic motivation, and developer experience in software engineering. In H. Sharp and T. Hall, editors, *Agile Processes in Software Engineering and Extreme Programming*, pages 104–117. Springer International Publishing, 2016. doi:10.1007/978-3-319-33515-5_9.

[52] C. F. Lam and S. T. Gurland. Self-determined work motivation predicts job outcomes, but what predicts self-determined work motivation? *Journal of Research in Personality*, 42(4):1109–1115, 2008. doi:10.1016/j.jrp.2008.02.002.

[53] F. J. Landy. *Psychology of work behavior*. Thomson Brooks/Cole Publishing Co, Belmont, CA, US, 4st edition, 1989.

[54] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015. doi:10.1038/nature14539.

[55] P. Lenberg, R. Feldt, and L. G. Wallgren. Behavioral software engineering: A definition and systematic literature review. *Journal of Systems and Software*, 107:15–37, 2015. doi:10.1016/j.jss.2015.04.084.

[56] J. Lerner and J. Tirole. Some simple economics of open-source. *Journal of Industrial Economics*, 50:197–234, Feb 2002. URL http://www.people.hbs.edu/jlerner/simple.pdf.

[57] Y. Li, C.-H. Tan, H. Teo, and T. Mattar. Motivating open source software developers: Influence of transformational and transactional leaderships. In *Proceedings of the 2006 ACM SIGMIS CPR Conference on Computer Personnel Research: Forty Four Years of Computer Personnel Research: Achievements, Challenges & the Future*, pages 34–43, Jan 2006. doi:10.1145/1125170.1125182.

[58] R. Likert. A technique for the measurement of attitudes. *Archives of Psychology*, 22:55, 1932.

[59] E. A. Locke. Toward a theory of task motivation and incentives. *Organizational Behavior and Human Performance*, 3(2):157–189, 1968. doi:10.1016/0030-5073(68)90004-4.

[60] C. Maslach, S. Jackson, and M. Leiter. Maslach burnout inventory third edition. In C. P. Zalaquett and R. J. Wood, editors, *Evaluating Stress: A Book of Resources*, pages 191–218. Scarecrow Education, 1997.

[61] A. H. Maslow. A theory of human motivation. *Psychological Review*, 50:370–396, 1943.

[62] J. D. Mayer, M. A. Faber, and X. Xu. Seventy-five years of motivation measures (1930–2005): A descriptive analysis. *Motivation and Emotion*, 31:83–103, 2007.

[63] D. C. McClelland. *The Achieving Society*. Van Nostrand, Princeton, NJ, 1961.

[64] A. N. Meyer, E. T. Barr, C. Bird, and T. Zimmermann. Today was a good day: The daily life of software developers. *IEEE Transactions on Software Engineering*, 47(5):863–880, 2021. doi:10.1109/TSE.2019.2904957.

[65] E. Murphy-Hill, C. Jaspan, C. Sadowski, D. Shepherd, M. Phillips, C. Winter, A. Knight, E. Smith, and M. Jorde. What predicts software developers' productivity? *IEEE Transactions on Software Engineering*, 47(3):582–594, 2021. doi:10.1109/TSE.2019.2900308.

[66] F. Murtagh and P. Contreras. Algorithms for hierarchical clustering: an overview. *WIREs Data Mining and Knowledge Discovery*, 2(1):86–97, 2012. doi:10.1002/widm.53.

[67] E. Oliveira, E. Fernandes, I. Steinmacher, M. Cristo, T. Conte, and A. Garcia. Code and commit metrics of developer productivity: a study on team leaders perceptions. *Empirical Software Engineering*, 25(4):2519–2549, Jul 2020. doi:10.1007/s10664-020-09820-z.

[68] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011. URL http://www.jmlr.org/papers/volume12/pedregosa11a/pedregosa11a.pdf.

[69] P. M. Podsakoff, S. B. MacKenzie, J.-Y. Lee, and N. P. Podsakoff. Common method biases in behavioral research: a critical review of the literature and recommended remedies. *Journal of Applied Psychology*, 88(5):879, 2003. doi:10.1037/0021-9010.88.5.879.

[70] J. R. Quinlan. *C4. 5: programs for machine learning*. Elsevier, 2014.

[71] S. Ramachandran and S. V. Rao. An effort towards identifying occupational culture among information systems professionals. In *Proceedings of the 2006 ACM SIGMIS CPR Conference on Computer Personnel Research: Forty Four Years of Computer Personnel Research: Achievements, Challenges & the Future*, pages 198–204, 2006. doi:10.1145/1125170.1125221.

[72] N. Raman, M. Cao, Y. Tsvetkov, C. Kästner, and B. Vasilescu. Stress and burnout in open source: Toward finding, understanding, and mitigating unhealthy interactions. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering: New Ideas and Emerging Results*, pages 57–60, 2020. doi:10.1145/3377816.3381732.

[73] A. J. Ratner, C. M. De Sa, S. Wu, D. Selsam, and C. Ré. Data programming: Creating large training sets, quickly. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 3567–3575. Curran Associates, Inc., 2016. URL http://papers.nips.cc/paper/6523-data-programming-creating-large-training-sets-quickly.pdf.

[74] E. S. Raymond. *The Cathedral and the Bazaar*. O'Reilly & Associates, Inc., Sebastopol, CA, USA, 1st edition, 1999.

[75] D. Riehle. The economic motivation of open source software: Stakeholder perspectives. *Computer*, 40(4):25–32, 2007. doi:10.1109/MC.2007.147.

[76] J. A. Roberts, I.-H. Hann, and S. A. Slaughter. Understanding the motivations, participation, and performance of open source software developers: A longitudinal study of the apache projects. *Manage. Sci.*, 52(7):984–999, Jul 2006. doi:10.1287/mnsc.1060.0554.

[77] G. Robles, L. Arjona Reina, A. Serebrenik, B. Vasilescu, and J. M. González-Barahona. Floss 2013: A survey dataset about free software contributors: Challenges for curating, sharing, and combining. In *Proceedings*

of the 11th Working Conference on Mining Software Repositories, pages 396–399, 2014. doi:10.1145/2597073.2597129.

[78] F. J. Roethlisberger and W. J. Dickson. *Management and the worker*. Harvard Univ. Press, Oxford, England, 1939.

[79] A. S. Tsui. A role set analysis of managerial reputation. *Academy of Management Proceedings*, 1982:265–269, Aug 1982. doi:10.5465/AMBPP.1982.4976642.

[80] J. Scott, E. O'Rourke, and A. Grant. Employee surveys are still one of the best ways to measure engagement. *Harvard Business Review*, March 2018. URL https://hbr.org/2018/03/employee-surveys-are-still-one-of-the-best-ways-to-measure-engagement.

[81] S. Scullen, M. Mount, and M. Goff. Understanding the latent structure of job performance ratings. *The Journal of Applied Psychology*, 85:956–70, Jan 2001. doi:10.1037//0021-9010.85.6.956.

[82] H. Sharp, N. Baddoo, S. Beecham, T. Hall, and H. Robinson. Models of motivation in software engineering. *Information & Software Technology*, 51(1):219–233, 2009. doi:10.1016/j.infsof.2008.05.009.

[83] H. Sharp and T. Hall. An initial investigation of software practitioners' motivation. In *Proceedings of the ICSE Workshop on Cooperative and Human Aspects on Software Engineering*, pages 84–91, 2009. doi:10.1109/CHASE.2009.5071418.

[84] H. Shavit and R. Shouval. Self-esteem and cognitive consistency effects on self-other evaluation. *Journal of Experimental Social Psychology*, 16(5):417–425, 1980. doi:10.1016/0022-1031(80)90048-7.

[85] J. A. Shepperd. Productivity loss in performance groups: A motivation analysis. *Psychological Bulletin*, 113(1):67, 1993. doi:10.1037/0033-2909.113.1.67.

[86] B. F. Skinner. *The Behavior of Organisms: An Experimental Analysis*. New York: Appleton-Century-Crofts, 1938.

[87] Stackoverflow. 2021 annual developer survey, 2021. URL https://insights.stackoverflow.com/survey/2021.

[88] R. Stallman. Why software should be free. 2007. URL https://www.gnu.org/philosophy/shouldbefree.en.html.

[89] F. R. Tanner. On motivating engineers. In *IEMC '03 Proceedings. Managing Technologically Driven Organizations: The Human Side of Innovation and Change*, pages 214–218, Nov 2003. doi:10.1109/IEMC.2003.1252263.

[90] M. Teghtsoonian and R. Teghtsoonian. How repeatable are stevens's power law exponents for individual subjects? *Perception & Psychophysics*, 10(3):147–149, 1971. doi:10.3758/BF03205774.

[91] J. Terrell, A. Kofink, J. Middleton, C. Rainear, E. Murphy-Hill, C. Parnin, and J. Stallings. Gender differences and bias in open source: Pull request acceptance of women versus men. *PeerJ Computer Science*, 3:e111, 2017. doi:10.7717/peerj-cs.111.

[92] V. Vapnik. *The Nature of Statistical Learning Theory*. Springer science & business media, 2013.

[93] V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16(2):264–280, 1971. doi:10.1137/1116025.

[94] P. F. Verhulst. Resherches mathematiques sur la loi d'accroissement de la population. *Nouveaux Memoires de l'Academie Royale des Sciences*, 18:1–41, 1845.

[95] C. V. Viswesvaran, D. Ones, and F. Schmidt. Comparative analysis of reliability of job performance ratings. *Journal of Applied Psychology*, 81:557–574, Oct 1996. doi:10.1037/0021-9010.81.5.557.

[96] G. von Krogh, S. Haefliger, S. Spaeth, and M. W. Wallin. Carrots and rainbows: Motivation and social practice in open source software development. *MIS Quarterly*, 36(2):649–676, 2012. URL http://www.jstor.org/stable/41703471.

[97] V. H. Vroom. *Work and Motivation*. Wiley, New York, 1964.

[98] J. Wen, S. Li, Z. Lin, Y. Hu, and C. Huang. Systematic literature review of machine learning based software development effort estimation models. *Information and Software Technology*, 54(1):41–59, 2012. doi:10.1016/j.infsof.2011.09.002.

[99] B. Wigert and J. Harter. Re-engineering performance management. *Gallup.com*, 2019. URL http://news.gallup.com/reports/208811/re-engineering-performance-management.aspx.

[100] T. A. Wright and R. Cropanzano. Psychological well-being and job satisfaction as predictors of job performance. *Journal of Occupational Health Psychology*, 5:84–94, 2000. doi:10.1037/1076-8998.5.1.84.

[101] Y. Ye and K. Kishida. Toward an understanding of the motivation open source software developers. In *Proceedings of the 25th International Conference on Software Engineering*, pages 419–429, 2003. doi:10.1109/ICSE.2003.1201220.

[102] B. Zadrozny and C. Elkan. Transforming classifier scores into accurate multiclass probability estimates. In *Proceedings of the 8th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 694–699, 2002. doi:10.1145/775047.775151.

A   Survey Questions

To facilitate the review, the full questionnaire used in the survey is reproduced herewith.

A.1   Survey introduction

Dear participant,

We are a team of researchers interested in improving software development (see for example https://www.cse.huji.ac.il/~feit/papers/Refactor19PROMISE.pdf).

If you contributed to a GitHub repository as a developer in the last 12 months, we ask for your help by answering questions about your contribution and motivation. Answering these questions is estimated to take 10–15 minutes of your time.

Based on the experience of respondents to this questionnaire in the past, you may gain new insights about your priorities in software development and areas of importance to you. Your answers, with the answers of others, will allow researchers in the future to investigate motivation, quality and productivity in software development and hopefully improve them.

We would appreciate a link to your GitHub profile in order to match your answers and GitHub activity (e.g., number of commits, years in the repository). We are aware that the profile is a personal identifier and we will keep it private and use it for research purposes only. The results of analysis of the profile data will be reported in aggregated form only. Of course, in case that you are not interested, you can leave the field empty.

If you are willing to participate in this study, move to the next page. By moving to the next page you agree to participate in this study. The only inconvenience that this study may cause you is the need to concentrate on the questions for about 10-15 minutes. Yet, you may quit this survey at any time without answering all the questions, with no consequences for you. We will be grateful if you complete ALL the questions. No personally identifying information will be collected, except your GitHub profile if you choose to share it.

Thank you so much for your help.

Prof. Dror Feitelson, Prof. Avi Kluger, Ph.D. candidate Idan Amit

If you have any question you can contact Idan Amit at idan.amit@mail.huji.ac.il

A.2 Questions regarding yourself

The questions in this section are in Likert scale where 1 is 'Strongly disagree' and 11 is 'Strongly agree'.

1. Productivity is more important to me than quality

2. My motivation has more influence on my productivity, than my skill

3. I regularly reach a high level of productivity (based on [65])

4. I am a relatively productive programmer

5. I am skilled in software development (based on [51])

6. My code is of high quality

7. I am satisfied with my performance in software development [51]

8. I want my code to be beautiful

9. I enjoy software development very much

10. It is important for me to program well (based on [51])

11. I write tests for my code

12. I write detailed commit messages

13. I contribute to open source in order to have an online portfolio

14. I try to write high quality code because others will see it

15. I enjoy trying to solve complex problems [4]

16. I contribute to open source in order to become a better programmer

17. I improved as a programmer since a year ago

18. I contribute to open source due to ideology

A.3 Questions regarding activity in a repository

Please choose *one specific* GitHub repository that you work on. Answer the following questions with respect to this repository. (These questions

- What is the link of the GitHub repository that you answer on? (Free ext)

- How many hours a week do you work on the repository (average)? (Free text)

- I'm being paid for my work in this repository (Yes/No)

The questions in this section are in Likert scale where 1 is 'Strongly disagree' and 11 is 'Strongly agree'.

1. I regularly have a high level of motivation to contribute to the repository (based on [65])

2. I have complete autonomy in contributing to the repository

3. I have significant influence on the repository

4. I feel responsible for the repository's success

5. I'm interested in the repository for my own needs

6. We have many heated arguments in the community. If you are the only developer in the project, please skip.

7. I wish that certain developers in the project will leave. If you are the only developer in the project, please skip.

8. My work on the repository is creative

9. Working on this repository is challenging

10. I derive satisfaction from working on this repository

11. The repository is important

12. When I look at what we accomplish, I feel a sense of pride.

13. Belonging to the community is motivating my work on the project. If you are the only developer in the project, please skip.

14. The community is very professional. If you are the only developer in the project, please skip.

15. I get recognition due to my contribution to the repository

16. I am a core member of the repository

17. I learn from my contributions

18. The quality of the code in this repository is better than others

19. Code quality in the repository improved since a year ago

20. The project's community of developers is more motivated than that of other projects. If you are the only developer in the project, please skip.

21. My personal motivation in this repository has increased since a year ago

22. In the past year, members of my project community put me down or were condescending to me. If you are the only developer in the project, please skip. (based on [26])

23. In the past year, members of my GitHub community made demeaning or derogatory remarks about me. If you are the only developer in the project, please skip. (based on [26])

24. In the past year, members of my project community asked questions that show their understanding of my contributions. If you are the only developer in the project, please skip. (based on [49])

25. In the past year, members of my project community expressed interest in my contributions. If you are the only developer in the project, please skip. (based on [49])

A.4   Job Satisfaction

The following questions are from Job Satisfaction Scale questionnaire [45]. We present the questionnaire as is in order to compare to previous results. In case that you find some questions irrelevant, please skip them.

The questions in this section are in Likert scale where 1 is 'Extremely dissatisfied' and 7 is 'Extremely satisfied', as in the original survey [45].

The questions indicate level of satisfaction with the following:

1. Freedom to choose your own method of working

2. Amount of variety in your work

3. Physical working conditions

4. Opportunities to use your abilities

5. Your colleagues and fellow workers

6. Recognition you get for good work

7. Your hours of work

8. Your remuneration (payment)

9. Amount of responsibility you are given

10. Taking everything into consideration, how do you feel about your work?

A.5   Demography

1. Country (Free text)

2. Age (0-100 selection)

3. Gender (Free text)

4. I work as a professional programmer (Yes/No)

5. Years of work experience (not including studies) (Free text)

6. Years of contribution to GitHub(0-15 selection)

7. Academic background (degree and graduation year) (Free text)

8. Git profile link (Free text) We would appreciate a link to your GitHub profile in order to match your answers and GitHub activity (e.g., number of commits, years in the repository). We are aware that the profile is a personal identifier and we will keep it private and use it for research purposes only. The results of analysis of the profile data will be reported in aggregated form only. Of course, in case that you are not interested, you can leave the field empty.

A.6   Open questions

1. Do you have any comments on the questionnaire or research? Are you motivated due to a cause that we didn't consider? Do you have a method that increases your code quality? (Free text)

2. Thank you for answering our survey. If you would like to be informed in the results of the research or to participate in the gift card lottery, please enter your email and we will send it to you once completed. The email will not be used for profile identification. (Free text)