3-6-2015

An Empirical Analysis of System-generated Data in Location-based Crowdsourcing

Carolin Durst

Michael Grottke

Follow this and additional works at: http://aisel.aisnet.org/wi2015

An Empirical Analysis of System-generated Data in Location-based Crowdsourcing

Carolin Durst[1], Michael Grottke[2]

[1] Department of Information Systems II, FAU Erlangen-Nürnberg, Germany
`Carolin.Durst@fau.de`
[2] Department of Statistics and Econometrics, FAU Erlangen-Nürnberg, Germany
`Michael.Grottke@fau.de`

Abstract. This paper develops a research model explaining how task location and incentives affect the take up and, for those tasks that are processed, the time to start. For an empirical analysis, we use the system-generated data of all 1860 location-based crowdsourcing tasks in Berlin available on the Streetspotr platform within one year.

The results indicate that while the population density of the task location does not influence the probability that some crowdworker will eventually process the task, a task located in a more densely-populated area tends to be taken up more quickly. Moreover, the take-up probability is expected to increase as the monetary and non-monetary incentives are raised. However, both increasing the monetary incentives and lowering the non-monetary incentives tends to shorten the time to start. This suggests that high non-monetary incentives with which unattractive tasks are endowed do not entice the crowdworkers to quickly set about processing these tasks.

Keywords: Location-based Crowdsourcing, Participation, Task Location, Incentives.

1   Introduction

The participatory generation of content (user-generated content) and the collective knowledge of a large number of users form the foundational pillars of Web 2.0. Companies apply the idea of the wisdom of the crowds [1] and collective intelligence to areas such as decision support, open innovation, social collaboration, and the so-called crowdsourcing [2]. The term "crowdsourcing", coined by Howe [3], refers to the outsourcing of traditional company tasks to the crowd, an indefinitely large and heterogeneous group of individuals. The crowd offers companies fast, flexible and relatively cost-efficient access to a large knowledge pool. The process is initiated with an open call on the Internet. Companies hope to obtain more efficient and better-quality results from the crowd, composed of experts and laypersons, than from an internally-developed solution [4, 5].

In conjunction with smartphones, crowdsourcing offers a new set of possibilities for the performance of crowdsourcing tasks. Due to their Internet functionality and

the integration of various sensors like motion sensors and GPS receivers, smartphones allow the users to generate and share content on the go, thus enabling simple and fast information exchange [6, 7]. With respect to the location of data collection, two types of mobile crowdsourcing can be distinguished: In location-independent (mobile) crowdsourcing, task solving is not linked to a specific location. An example for this kind of crowdsourcing is the mobile application of Amazon's Mechanical Turk, which enables crowdworkers to work on tasks location-independently using their mobile phones [8, 9]. In contrast, location-based crowdsourcing requires the presence of a crowdworker at a specific location, because a certain activity (e. g., the collection of location-dependent information) can only be conducted on-site [10]. In both cases, the results can be submitted either via a mobile device or via a desktop PC [11].

Although crowdsourcing is gaining widespread popularity, location-based crowdsourcing has only received little attention in the literature. However, there is a need for better insights into this specific form of crowdsourcing: Companies involved in location-based crowdsourcing projects are faced with new challenges due to the geographical constraints of the tasks or the limited screen size of mobile phones. If companies understand the effects of different task design parameters on the performance of the mobile workforce, location-based crowdsourcing holds a considerable potential, due to its aforementioned opportunities. Baily and Fessler [12] argued that if a task is simple and relatively unattractive for the crowd, higher monetary incentives are likely to lead to higher take-up rates, and improve the overall performance. Regarding the concept of location, researchers have often defined location as the distance between the location of a specific crowdworker and the location of the task, and have found that crowdworkers prefer to receive location-based tasks in close proximity to their home or their reporting location [11, 13]. These researchers thus considered "location" as a concept combining individual and task attributes. In our research paper, we refer to the "task location" as a characteristic related to the task alone, independent of any crowdworker's location. Attributes related to the task location are the geographic coordinates and the population density of the task location, for instance.

Using this task location concept, we empirically investigate the effects of task location, monetary incentives and non-monetary incentives on participation. To this end, we make use of system-generated data from the crowdsourcing platform Streetspotr. First, we analyze how the identified parameters influence the take-up probability. Second, we take a closer look at the time to start, which refers to the time elapsed between the moment a task is issued on the crowdsourcing platform and the moment any crowdworker starts processing it [14]. To the best of our knowledge, there is no prior research explicitly studying these relationships using actual performance data. Insights into these relationships can help researchers and practitioners to design tasks in a way ensuring that they are indeed accepted up by crowdworkers, and more quickly at that.

The remainder of the paper is organized as follows. After elaborating on job design theories, we review those crowdsourcing studies investigating the influence of location and incentives on participation. We then develop our research model and apply it to system-generated data from the crowdsourcing platform Streetspotr. Finally, we discuss our findings and the limitations of the study.

2   Theoretical Background

Our research model is rooted in job design theory, which states that organizational and individual needs can both be met effectively through the manipulation of certain job characteristics [15]. Torraco [16] conducted a comprehensive literature review on job design theories and their application to new work environments. To study job design on the individual task level, he considered using the job characteristics model [17]. This model systematically outlines the links between the characteristics of a job, the individual's experience with these job characteristics, and the outcomes in terms of motivation, satisfaction and performance. Hackman and Oldham [17] identified five core job dimensions: skill variety, task identity, task significance, autonomy, and feedback. The job characteristics model is among the most-recognized and complete theories for explaining job design characteristics and their relationships to work performance and motivation, and it has led to an impressive body of research on job design. However, Torraco [16] found that this model and the related research focus on traditional work environments, and that they thus fail to explain the effects of job design in new work environments, where different task design parameters are of interest. He argued that activity theory – also belonging to the family of job design theories – may instead turn out to be capable of explaining job design in future settings.

Activity theory stems from the Soviet cultural-historical psychology of Vygotsky, Leont'ev, and Luria [18]. It facilitates the analysis of purposeful behavior by focusing on the structure of the activity itself. In Vygotsky's model of mediated act [19] the subject of any activity is an individual who is engaged in the activity. Mediated by tools, the individual transforms the object of the activity and achieves a certain outcome (see **Fig. 1**). In this context, an activity is a goal-directed interaction of a subject with an object through the use of tools.


Fig. 1. Vygotsky's model of mediated act

Activity theory is considered to be a meta-theory or a framework, rather than a predictive theory. The considerable flexibility of this theory resulted in a number of applications in different areas. Engeström [20] used it to study the redesign of work in a pediatric health care facility, whereas Jonassen and Rohrer-Murphy [21] as well as Karasavvidis [22] proposed activity theory as a framework to design learning environments. Kuutti [18] and Nardi [23] suggested that it can serve as an alternative framework for studying human-computer interaction and design. In his work "Activity theory as a theoretical foundation for information systems research", Ditsa [24] also

described the potential of activity theory to combine human and technological aspects of information systems in a more holistic research approach. Until today, many researchers in the field of information systems management and human-computer interaction have applied activity theory (e. g., [25–29]). Recently, Hautasaari [30] used activity theory in the context of crowdsourcing and analyzed information search and translation activities. The results of his activity analysis were used to define design implications for a Wikipedia translation support system. Asmolov [31] argued that crowdsourcing platforms can give rise to different types of new activity systems and that activity theory can assist in conceptualizing the relationships between subject and object, as well as in analyzing the relationships around crowdsourcing platforms.

In our research project, we also draw on activity theory to investigate the performance of a crowdworker who has taken up a task, for example in terms of the time he needs to complete the task. According to activity theory, this time to completion is influenced by aspects of the object (i. e., task design parameters like its severity), characteristics of the subject (e. g., the crowdworker's age and education), as well as properties of the tools used to fulfill the task (e. g., the quality of the camera built into the crowdworker's smartphone). However, in this paper we will study that period in the task life cycle before any individual has begun to work on it. The outcomes of this period in terms of the take up (i. e., the fact whether or not a task is processed by any crowdworker) and, for those tasks that are taken up, the time to start (i. e., the time elapsed from issuing the task until it is taken up by some crowdworker) are therefore related to the design parameters of the task itself, which refer to the characteristics of the object. For the analysis of this period in the task life cycle, the outcomes are not directly related to the characteristics of any specific crowdworker or any specific tool. This is especially clear when it comes to explaining why a certain task has not been taken up; for such a task, an association with a specific crowdworker or a specific tool has never been established. Even for those tasks that have been processed, the time elapsed before the crowdworker finally took it up is not only related to the characteristics of this one crowdworker and his or her smartphone. Rather, it is a consequence of the behavior of the whole set of crowdworkers who could potentially have taken up the task. Statistically speaking, the observed time to start is the first order statistic of the reaction times of all these individuals, who are representing a set of competing risks.

3   Related Work

A number of studies have already shed light on the relationships between different crowdsourcing parameters and participation. Many well-cited studies (e. g., [32–34]) focus on micro tasks outsourced to Amazon's Mechanical Turk platform. However, only very few studies (e. g., [11, 13, 35]) have been conducted in the context of location-based crowdsourcing. Therefore, this section reviews all kinds of crowdsourcing studies investigating the relationships between location, incentives and participation. In the next section, we will draw on the investigated variables and the findings of these studies when developing our research model and the related hypotheses.

3.1 Location and Participation

It seems obvious that in location-based crowdsourcing the location of a task should have a bearing on participation. However, there has been surprisingly little work examining the relationship between location and participation in location-based crowdsourcing. To study the effect of location on participation, Alt et al. [11] implemented a crowdsourcing platform that integrated location as a parameter for allocating tasks to crowdworkers. Conducting two user studies, they found that participants prefer to solve tasks at home or in its vicinity and that participants prefer to search for tasks in their current location. These two user studies employed a small sample of nine participants, and they were of an experimental and qualitative nature. Following a similar qualitative research approach, Väätäjä et al. [13] studied location-based crowdsourcing for news reporting. To develop a new crowdsourcing process in hyper-local news production, they recruited nine participants for their first study and 19 participants for the subsequent quasi-experiment in field conditions. Their results indicate that the location affects the willingness to receive location-based crowdsourcing assignments. In both research projects, the location-based tasks could be accessed and accomplished by a crowdworker only via smartphone and only if the participant was at the specific location. Thus the distance between the crowdworker and the crowdsourcing task played an important role in the respective context.

In other flavors of location-based crowdsourcing, location is certainly influential but not crucial for participation. One main research stream in location-based crowdsourcing refers to participatory sensing, in which location-sensitive data is collected with mobile sensors [36]. In the project BrusSense, for instance, participants use their mobile phones to record noise levels at different locations throughout a city. The aggregated homogeneous contributions result in a Noise Exposure Map to monitor the noise pollution in the city [37]. While the information is linked to a specific location, each participant can perform the task at any location within a given geographic area.

Literature thus contains evidence for a relationship between location and participation in certain settings, whereby in most of the cases location is defined as the distance between the location of the crowdworker and the location of the task. The concept "distance" is therefore inseparably linked with the individual crowdworker. As discussed above, our study focuses on that time period of the task life cycle in which the design parameters of the task play a pivotal role for the outcomes. Therefore, we will include task location in terms of population density of the location of the task in our research model.

3.2 Incentives and Participation

Mason and Watts [33] studied an image ordering task and a word puzzle task to investigate how compensation affects performance on Amazon's Mechanical Turk platform. They found that higher financial compensation increases the quantity of work performed by the participants, while it does not necessarily improve its quality. Using Amazon's Mechanical Turk to collect questionnaire data for research in psychology and in social sciences, Buhrmester et al. [34] showed that the participation rate is

affected by the compensation rate and the length of the task. Bailey and Fessler [12] investigated moderating effects of task complexity and task attractiveness on the impact of monetary incentives and found that higher compensation rates improve participation if the task is simple. In full agreement with the previous findings, Rogstadius et al. [38] showed that rewards substantially increase both the take-up and the overall completion rates. Monetary incentives were also found to be the most important factor for participation in the software development crowdsourcing domain [39]. Zheng et al. [40] investigated the relationships between crowdsourcing contest characteristics, motivation, and participation. In contrast to other studies, their results did not support the hypothesis that monetary incentives positively influence participation; however, they found that the possibility of gaining recognition was positively associated with participation. Paolacci et al. [41] analyzed survey tasks on Amazon's Mechanical Turk as an alternative to online surveys. They found that there are many non-monetary reasons for participation, such as entertainment or simply "killing time".

A great number of studies addressing the relationship between incentives and participation (e. g., [33, 34, 42]) indicate the relevance of this topic in research. The majority of the presented studies suggest that the amount of incentives influences participation. We will therefore include both monetary and non-monetary incentives in our research model.

4 Research Model

The above discussion of the related work suggests that both the task location and the incentives may influence participation in location-based crowdsourcing. Apart from qualitative or experimental work, there has not been any comprehensive study that empirically investigates the relationships between task location, monetary and non-monetary incentives as well as participation in a location-based crowdsourcing setting and using actual data from a crowdsourcing platform; specifically, the trade-off effects between task location and incentives have not yet been analyzed. If we understand these relationships, we can provide companies with design guidelines for their crowdsourcing tasks, to improve the outcomes in the first period of the task life cycle. Tailored to the specific circumstances of location-based crowdsourcing, we propose the research model shown in **Fig. 2**, which is based on activity theory and Vygotsky's model of mediated act [19]. It incorporates the previously-identified independent variables task location, monetary incentives and non-monetary incentives as well as the dependent variables take up and time to start.
The corresponding hypotheses are as follows:

H1: Task location is associated with the take-up probability.
H2: Monetary incentives are positively associated with the take-up probability.
H3: Non-monetary incentives are positively associated with the take-up probability.
H4: Task location is associated with the time to start.
H5: Monetary incentives are positively associated with the time to start.
H6: Non-monetary incentives are positively associated with the time to start.

Fig. 2. Research Model

5 Research Method and Data Analysis

5.1 Research Setting and Data Collection

The data for this study comes from the crowdsourcing platform Streetspotr. The Streetspotr platform serves as a service intermediary and enables companies to outsource small location-based tasks to private individuals. A company can configure tasks on the crowdsourcing platform, which are subsequently made available and which can then be processed with a mobile application. To create a task a company must decide on the task title, the task description, the task location, and the incentives. The company then needs to choose a suitable tool set configuration using one or several of the tool set options text, rating, photo, video, single choice and multiple choice provided by the system. A typical task in the field of retail execution could be: "Take a picture of the 'Death Star' (Lego number 10143) and its presentation (decoration and display) in the Lego Store. Ensure that the product and its environment can easily be seen." Registered crowdworkers are notified about the new assignment via an open call; they can review, accept and process the job from their smartphones. Once a task has been completed, the company is informed through the crowdsourcing platform and reviews the result. If the result is satisfactory, the crowdworker receives the dedicated incentive in the form of Euros and/or virtual points (so-called Streetpoints). By collecting Streetpoints, the crowdworker gains recognition via public leaderboards within the crowdworker community. To date, Streetspotr counts over 200,000 registered users in Germany, 72% of which are male, while only 28% are female.

We carried out our analyses employing all user-activity data from the Streetspotr platform generated in Berlin between April 1, 2012 and May 31, 2013. For each one of the 1860 unique location-based crowdsourcing tasks in this sample, we extracted the available information on task location, incentives, and participation from the Streetspotr database via SQL statements. **Table 1** lists descriptive statistics for all variables extracted, including measures of location (mean, median), measures of variation (standard deviation), and the number of tasks in our sample for which the value of the respective variable was available.

For our purposes, a crucial aspect of the task location seems to be if the task can easily and quickly be reached by the crowdworkers, or if it is located in some far-off region. We therefore operationalized task location by the population density of the district (Ortsteil) in which the task is located.

Table 1. Descriptive statistics for the variables extracted

	Minimum	Maximum	Mean	Median	Std. deviation	Sample size
PopDensity	0.0397	1.6164	0.8849	1.0713	0.4061	1860
IncEuro	0	15	1.0435	1	1.4844	1860
IncPoints	0	120	31.6290	15	19.4147	1860
TakeUp	0	1	0.9909	NA	NA	1860
TTS	0.0028	453.3925	107.2566	13.9386	134.0516	1843

Berlin counts 96 districts with population densities ranging from 150 inhabitants per square kilometer in Blankenfelde to 16,261 inhabitants per square kilometer in Friedenau. From the Streetspotr database we extracted the geographic coordinates of each task and performed a reverse geocoding to convert the coordinates into addresses including the district. After the automated reverse geocoding process, we ran consistency checks on the results. In some cases the district information was missing or the address did not match the respective district. After identifying all inconsistencies, we manually added missing and corrected false information. Finally, we mapped the district density (PopDensity), measured in 10,000 inhabitants/km², available from official sources [43] to the converted addresses. As can be seen from **Table 1**, the 1860 tasks were located in districts with population densities between 397 and 16,164 inhabitants per square kilometer.

Monetary incentives are measured in Euro (IncEuro), and non-monetary incentives are operationalized by the Streetpoints (IncPoints) a crowdworker receives when he or she completes the task. It was possible to directly extract both measures from the database. While IncEuro ranged from zero to 15 Euros, less than 50 per cent of all tasks were endowed with monetary incentives exceeding one Euro, such that the mean of IncEuro is only slightly higher than one Euro. In contrast, the mean and the median of IncPoints are not that close to their minimum value, indicating less concentration at the lower end of the range of Streetpoints promised.

Extracting the times of all status changes for a task documented in the database, we were able to derive a binary variable (TakeUp) indicating whether the task was *at all* taken up by any crowdworker (TakeUp = 1) or whether no crowdworker has ever started to process it (TakeUp = 0). While a task might be published again after being unsatisfactorily completed by a crowdworker, we did not take into account the crowdworkers' reaction to such a repeated advertisement, which might differ from the reaction to a new task. Hence, TakeUp takes exactly one value for each of the 1860 tasks in our sample. Although TakeUp is a qualitative variable, its mean has a useful interpretation; it represents the fraction of tasks taken up by a crowdworker: 99.09% (i. e., 1843) of the total 1860 tasks were ever started to be processed.

For each task with TakeUp = 1, we also calculated the time to start (TTS) as the time interval (in days) between the moment when the task was *originally* made available and the moment when it was *first* taken up by a crowdworker. Overall, 664 unique crowdworkers initially worked on any of the 1843 tasks that were taken up.

Their gender distribution (23% female, 76% male, gender not provided by 1%) reflects the one of the Streetspotr population.

5.2 Analysis

All analyses described in this paper were carried out with the statistical software package R [44]. To examine whether there is an association between our three independent variables and the fact that a task is (not) taken up by any crowdworker, as suggested by hypotheses H1 to H3, we employ a logit model of the form

$$\ln\left(\frac{P(\text{TakeUp} = 1)}{1 - P(\text{TakeUp} = 1)}\right) = \beta_{10} + \beta_{11} \cdot \text{PopDensity} + \beta_{12} \cdot \text{IncEuro} + \beta_{13} \cdot \text{IncPoints} + \varepsilon_1. \qquad (1)$$

The variance inflation factors calculated for the independent variables based on the 1860 observed tasks amount to 1.0013 (PopDensity), 1.0015 (IncEuro) and 1.0003 (IncPoints), respectively. All these values are close to one. This means that for each independent variable the variance of the estimator of the related parameter is hardly higher than in the hypothetic case in which there is no linear dependence between all independent variables [45]. Multicollinearity is thus not an issue in our data set.

Table 2 lists the parameter estimates obtained via iteratively reweighted least squares estimation, their standard errors, as well as the related p values and odds ratios. All quantities have been rounded to four decimal places; p values below 0.00005 are therefore shown as 0.0000. According to these results, there is no significant association between the population density and the probability that the task will be taken up by a crowdworker, which contradicts our hypothesis H1. However, we do find support for hypotheses H2 and H3. At any reasonable significance level, the null hypothesis that IncEuro (or IncPoints) does not influence the participation probability can be rejected, and the estimated associations are positive: The higher the monetary or non-monetary incentives, the higher the probability that some crowdworker will start working on the task. More precisely, each odds ratio shown represents the factor by which the take-up odds (i. e., the probability of the task being taken up divided by the probability of not being taken up) tend to change if the value of the respective explanatory variable is increased by one. For example, if one Euro more has been promised as an incentive for task A than for task B (everything else being the same), then task A is 40.80 more likely, than not, to be taken up as compared with task B. It can thus be seen that increasing the monetary incentives by one Euro has a much larger effect than increasing the non-monetary incentives by one Streetpoint.

Table 2. Regression results for the logit model in Equation (1)

|	Estimate	Standard error	p value	Odds ratio, exp(Estimate)
(Intercept)	-1.1115	0.9274	0.0000	NA
PopDensity	-0.0433	0.7148	0.9520	0.9576
IncEuro	3.7088	0.5654	0.0000	40.8042
IncPoints	0.1266	0.0177	0.0000	1.1350
LR: 62.9701 (df: 3, p value: 0.0000)				

Fig. 3. Take-up probability estimated for each task based on the logit model

The overall fit of the logit model is evaluated in terms of the likelihood ratio (LR) statistic, which compares the likelihood value attained by this model with the one of the intercept-only model. Since the LR statistic asymptotically follows a chi-square distribution with three degrees of freedom (df), the value of 62.9701 attained implies a p value below 0.00005. We can therefore reject the null hypothesis that the intercept-only model is as effective as our model at any reasonable significance level.

Fig. 3 depicts the take-up probability estimated using the logit model for each of the 1860 tasks (ordered based on their system-internal task ID). While those 17 tasks in the data set that have never been taken up by a crowdworker are drawn as red triangles, the 1843 tasks for which TakeUp = 1 are represented by black circles. Indeed, the estimated take-up probabilities seem to allow a good separation between the two types of tasks. For 14 of the 17 unclaimed tasks the estimated probability ranges between 0.67 and 0.69. Only four of the tasks processed by a crowdworker, namely tasks no. 1137, 1138, 1139, and 1141 in the data set, feature an estimated take-up probability below 0.9. In fact, these are the only tasks observed for which neither any monetary nor any non-monetary incentives had been promised.

According to the hypotheses H4 to H6 the variables PopDensity, IncEuro and IncPoints should also be associated with TTS, the time until the task is taken up by a crowdworker. Using TTS as the dependent variable in a linear regression model is problematic, because TTS is restricted to non-negative values, while the model might produce fitted values below zero. We therefore propose the following regression model explaining the natural logarithm of TTS:

$$\ln \text{TTS} = \beta_{20} + \beta_{21} \cdot \text{PopDensity} + \beta_{22} \cdot \text{IncEuro} + \beta_{23} \cdot \text{IncPoints} + \varepsilon_2. \qquad (2)$$

This model can be estimated based on those 1843 tasks that have been processed. As shown before, for the full data set of all 1860 tasks there is no multicollinearity between the three independent variables, and this conclusion still holds after dropping the few tasks for which TakeUp = 0.

In the classic linear regression model, it is assumed that the disturbances ε_2 follow a normal distribution. This assumption ensures that the t tests employed for judging the significance of associations are valid. We use the Shapiro-Wilk test [46] for determining whether or not the residuals, i. e., the estimated disturbances, have been sampled from a normal distribution. The value of the test statistic obtained from our data (0.9168) implies a p value smaller than 0.00005; for any reasonable significance level, the hypothesis of normally-distributed disturbances thus needs to be rejected. While this means that the t test statistics do not exactly follow a t distribution, it can be shown that the t tests are asymptotically valid, and are thus approximately valid for large sample sizes, if the other assumptions of the classic linear model (known as the Gauss-Markov assumptions) hold [47].

Our large number of observations would surely allow us to make use of asymptotic results. However, the assumption that the disturbances are homoscedastic, which is one of the Gauss-Markov assumptions, seems questionable. To check it, we carry out a Breusch-Pagan test [48], in the studentized version due to Koenker (1981). Indeed, the test results (BP test statistic = 206.736, p value < 0.00005) indicate that we need to reject the hypothesis that the homoscedasticity assumption holds. As a consequence, the usual equation for calculating the standard errors of the parameter estimators are not valid [47], and the t test statistics derived based on these standard errors are not (asymptotically) normal.

White [50] proposed an approach for computing valid standard errors in the presence of heteroscedasticity, and we make use of these heteroscedasticity-consistent (HC) standard errors. They are shown in **Table 3**, listing our regression results for the linear model, in addition to the related p values and the parameter estimates. Obviously, at any reasonable significance level the hypothesis that the independent variable does not affect the logarithm of the TTS can be rejected for PopDensity, IncEuro, and IncPoints. The association between IncPoints and the logarithm of TTS is positive, which is counterintuitive. In Section 6, we will give a rationale for this finding.

However, for the former two explanatory variables, there is a negative association, as expected: Tasks located in a district with a higher population density tend to be taken up more quickly by a crowdworker, as are tasks for which a higher monetary incentive has been promised. For example, if the population density of the location of task A exceeds the one of task B by of 5,000 inhabitants per square kilometer, then it can be expected that the time to start of task A will be 15.3% (exp(-0.3330·0.5)-1) lower than the one of task B. By increasing the monetary incentives by 16 Cents a similar decrease in the time to start of 15.0% (exp(-1.0140·0.16)-1) could be attained.

Table 3. Regression results for the linear model in Equation (2)

	Estimate	HC standard error	p value
(Intercept)	1.3410	0.1479	0.0000
PopDensity	-0.3330	0.0913	0.0003
IncEuro	-1.0140	0.0435	0.0000
IncPoints	0.0851	0.0023	0.0000
R^2: 0.6975			

The coefficient of determination (R^2) obtained for the linear regression model indicates that 69.75% of the total variation in the logarithm of TTS can be explained by the three independent variables, representing a very satisfactory model fit.

6   Discussion

Our results show support for different findings from prior literature. We can confirm that monetary and non-monetary incentives both positively influence participation [12, 34, 38, 39, 41], as far as take up is concerned. However, monetary incentives show a much larger effect on take up than the non-monetary Streetpoints. Although a crowdworker can gain recognition within the Streetspotr community by collecting Streetpoints, money is still a more powerful incentive. Even if we could not confirm the hypothesis that the task location affects the probability that some crowdworker will start working on the task, we found that it influences participation in terms of the time to start. In many use cases of location-based crowdsourcing the location of a task is a fixed parameter that cannot be changed by the company designing the task (e. g., for the task of taking a photo of a specific building). However, we found that the expected effects of large differences in population density correspond to the ones of rather small changes in the monetary incentives. Thus, monetary incentives can compensate for "unattractive locations". In contrast to Zheng et al. [40] our results indicate that non-monetary incentives increase the time to start. Why should crowdworkers tend to be more hesitant to start working on tasks for which they will receive more Streetpoints? The reason for this finding seems to be rooted in the fact that tasks perceived to be unattractive by the company (either due to their location, or because the company is not willing to promise a substantial Euro amount) tend to be endowed with higher non-monetary incentives. While we have seen before that increasing the number of Streetpoints increases the probability of taking up the task, it does not entice the crowdworkers to do so more quickly. Even in the presence of a high non-monetary incentive, an unattractive task tends to be taken up and processed by a crowdworker when it is convenient for him or her (for example, because (s)he happens to pass by in its vicinity); by themselves, the Streetpoints do not seem to allure a crowdworker to process the task at his or her earliest convenience.

These findings could be used to implement an automated task recommendation system for companies in the backend of the crowdsourcing platform. For example, such a system could recommend how to set the monetary incentives if a company wants to make sure that the task can be expected to be taken up within a certain period of time. Moreover, our results hint on how to improve the general design of location-based crowdsourcing platforms: As monetary and non-monetary incentives have been found to significantly influence the take-up behavior of the crowdworkers, it may be a good idea to prominently present this information for each task within the mobile application. Of course, it should be noted that exploiting the detected relationships in these ways might change the effects in the future. For instance, if a redesign of the graphical user interface should give more visibility to the Euro amount promised for finishing a task, this might make more crowdworkers focus on the monetary incen-

tives when selecting a task to process, leading to an even stronger positive association between this variable and the take-up probability.

The findings of this study need to be weighed against its potential limitations. Our operationalization of task location by the population density of the district in which the task is located might pose a threat to construct validity. One could argue that the number of crowdworkers present within a certain radius from the task might be a better metric. However, movement profiles of the crowdworkers are not available, and the registered home address of a crowdworker does not necessarily represent the location where (s)he spends most of his or her time. Also, using such information may be problematic for reasons of data privacy. We therefore think that the chosen operationalization is the best currently available. Furthermore, as we have constructed our research model from previous literature, and have used appropriate statistical techniques taking into account properties like non-normal and heteroscedastic disturbances, we assume that the internal validity of our conclusions is not at risk. As for external validity, the data analyzed in this paper represents the location-based crowdsourcing tasks in Berlin in about one year's time. While for example the population density varies widely even within this one city, we cannot be sure that the results of this study also apply to rural areas. Moreover, it is possible that the behavior of the crowdworkers and their reaction to incentives will change in the future.

In future studies, we will investigate additional task design parameters as well as the amount of time spent on completing a task. Both a fast acceptance and a rapid processing time ensure that companies get their crowdsourcing results quickly.

References

1. Surowiecki, J.: The Wisdom of Crowds. Anchor Books, New York (2005)
2. Leimeister, J.M.: Collective Intelligence. Bus. Inf. Syst. Eng. 2, 245–248 (2010)
3. Howe, J.: The Rise of Crowdsourcing. Wired Mag. 14, 1–5 (2006)
4. Brabham, D.C.: Crowdsourcing as a Model for Problem Solving: An Introduction and Cases. Converg. Int. J. Res. into New Media Technol. 14, 75–90 (2008)
5. Hammon, L., Hippner, H.: Crowdsourcing. Wirtschaftsinformatik. 54, 165–168 (2012)
6. Oulasvirta, A., Rattenbury, T., Ma, L., Raita, E.: Habits Make Smartphone Use More Pervasive. Pers. Ubiquitous Comput. 16, 105-114 (2012)
7. Oinas-Kukkonen, H.: Behavior Change Support Systems: A Research Model and Agenda. In: 5th Int. Conf. on Persuasive Technology, pp. 4–14 (2010)
8. Eagle, N.: txteagle : Mobile Crowdsourcing. In: Int. Conf. on Internationalization Design and Global Development, pp. 447–456 (2009)
9. Yan, T., Marzilli, M., Holmes, R., Ganesan, D., Corner, M.: mCrowd – A Platform for Mobile Crowdsourcing. In: 7th ACM Conf. on Embedded Networked Sensor Systems, pp. 347-348 (2009)
10. Väätäjä, H., Vainio, T., Sirkkunen, E., Salo, K.: Crowdsourced News Reporting: Supporting News Content Creation with Mobile Phones. In: 13th Int. Conf. on Human Computer Interaction with Mobile Devices and Services, pp. 435–444 (2011)
11. Alt, F., Shirazi, A.S., Schmidt, A., Kramer, U., Nawaz, Z.: Location-based Crowdsourcing: Extending Crowdsourcing to the Real World. In: 6th Nordic Conf. on Human-Computer Interaction: Extending Boundaries, pp. 13-22 (2010)

12. Bailey, C.D., Fessler, N.J.: The Moderating Effects of Task Complexity and Task Attractiveness on the Impact of Monetary Incentives in Repeated Tasks. J. Manag. Account. Res. 23, 189–210 (2011)
13. Väätäjä, H., Vainio, T., Sirkkunen, E.: Location-based Crowdsourcing of Hyperlocal News: Dimensions of Participation Preferences. In: 13th Int. Conf. on Human Computer Interaction with Mobile Devices and Services, pp. 435-444 (2012)
14. Junglas, I., Abraham, C., Watson, R.T.: Task-Technology Fit for Mobile Locatable Information Systems. Decis. Support Syst. 45, 1046–1057 (2008)
15. Buchanan, D.A.: The Development of Job Design Theories and Techniques. Saxon House (1979)
16. Torraco, R.J.: Work Design Theory: A Review and Critique with Implications for Human Resource Development. Hum. Resour. Dev. Q. 16, 85–109 (2005)
17. Hackman, J.R., Oldham, G.R.: Motivation through the Design of Work: Test of a Theory. Organ. Behav. Hum. Perform. 16, 250–279 (1976)
18. Kuutti, K.: Activity Theory as a Potential Framework for Human-Computer Interaction research. In: Nardi, B.A. (ed.) Context and Consciousness: Activity Theory and HCI, pp. 17-44 (1996)
19. Vygotsky, L.S.: Mind in Society: The Development of Higher Psychological Processes. Harvard University Press, Cambridge (1978)
20. Engeström, Y.: Activity Theory as a Framework for Analyzing and Redesigning Work. Ergonomics. 43, 960–74 (2000)
21. Jonassen, D., Rohrer-Murphy, L.: Activity Theory as a Framework for Designing Constructivist Learning Environments. Educ. Technol. Res. 47, 61-79 (1999)
22. Karasavvidis, I.: Understanding Wikibook-Based Tensions in Higher Education: an Activity Theory Approach. E-Learning Digit. Media. 7, 386-394 (2010)
23. Nardi, B.: Activity Theory and Human-Computer Interaction In: Nardi, B.A. (ed.) Context and Consciousness: Activity Theory and Human-Computer Interaction, pp. 4-8 (1996)
24. Ditsa, G.: Activity Theory as a Theoretical Foundation for Information Systems Research. In: Ditsa. G.: Information Management: Support Systems and Multimedia Technology, pp. 192-231 (2003)
25. O'Leary, D.E.: An Activity Theory Framework for DSS for Extreme Events. In: 2010 Conf. on Bridging the Socio-technical Gap in DSS: Challenges for the Next Decade, pp. 487-497 (2010)
26. Hamid, S., Waycott, J., Kurnia, S.: The Use of Online Social Networking for Higher Education from an Activity Theory Perspective. In: 14th Pacific Asia Conf. on Inf. Systems, pp. 1414-1425 (2010)
27. Kamaruddin, K.A., Yusop, N.S.M., Ali, M.A.M.: Using Activity Theory in Analyzing Requirements for Mobile Phone Application. In: 5th Malaysian Conf. in Software Engineering, pp. 7–13. (2011)
28. Bharosa, N., Lee, J., Janssen, M., Rao, H.R.: An Activity Theory Analysis of Boundary Objects in Cross-Border Inf. Systems Development for Disaster Management. Secur. Inform. 1, 1-17 (2012)
29. Barki, H., Titah, R., Boffo, C.: Information System Use-Related Activity: An Expanded Behavioral Conceptualization of Individual-Level Information System Use. Inf. Syst. Res. 18, 173–192 (2007)
30. Hautasaari, A.: "Could Someone Please Translate This?" - Activity Analysis of Wikipedia Article Translation by Non-Experts. In: Conf. on Computer Supported Cooperative Work, pp. 945–954 (2013)

31. Asmolov, G.: Crowdsourcing as an Activity System: Online Platforms as Mediating Artifacts - A Conceptual Framework for the Comparative Analysis of Crowdsourcing in Emergencies. In: Sintelnet WG5 Workshop on Crowd Intelligence: Foundations, Methods and Practices, pp. 24–42 (2014)
32. Horton, J.J., Chilton, L.B.: The Labor Economics of Paid Crowdsourcing. In: 11th ACM Conf. on Electronic Commerce, pp. 209–218 (2010)
33. Mason, W., Watts, D.J.: Financial Incentives and the "Performance of Crowds". In: ACM SIGKDD Workshop on Human Computation, pp. 100–108 (2009)
34. Buhrmester, M., Kwang, T., Gosling, S.D.: Amazon's Mechanical Turk: A New Source of Inexpensive, Yet High-Quality, Data? Perspect. Psychol. Sci. 6, 3–5 (2011)
35. Väätäjä, H., Egglestone, P.: Briefing News Reporting with Mobile Assignments: Perceptions, Needs and Challenges. In: ACM 2012 Conf. on Computer Supported Cooperative Work, pp. 485–494 (2012)
36. Thebault-Spieker, J.: Crowdsourced Participatory Sensing : Applications and Motivation of Work, http://www-users.cselabs.umn.edu/classes/Spring-2012/csci8002/PAPERS/jacob.pdf (Accessed: 14.08.2014)
37. D'Hondt, E., Stevens, M.: Empowering Technologies for Environmental Participatory Sensing, http://www.brussense.be/Pubs/files/empowering-technologies-for-environmental-participatory-sensing.pdf (Accessed: 22.03.2013)
38. Rogstadius, J., Kostakos, V., Kittur, A., Smus, B., Laredo, J., Vukovic, M.: An Assessment of Intrinsic and Extrinsic Motivation on Task Performance in Crowdsourcing Markets. In: 5th Int. AAAI Conf. on Weblogs and Social Media, pp. 321-328 (2011)
39. Pilz, D., Gewald, H.: Does Money Matter? Motivational Factors for Participation in Paid- and Non-Profit-Crowdsourcing Communities. In: 11th Int. Conf. on Wirtschaftsinformatik, pp. 577–591 (2013)
40. Zheng, H., Li, D., Hou, W.: Task Design, Motivation, and Participation in Crowdsourcing Contests. Int. J. Electron. Commer. 15, 57–88 (2011)
41. Paolacci, G., Chandler, J., Stern, L.N.: Running Experiments on Amazon Mechanical Turk. Judgm. Decis. Mak. 5, 411–419 (2010)
42. Bonner, S.E., Sprinkle, G.B.: The Effects of Monetary Incentives on Effort and Task Performance: Theories, Evidence, and a Framework for Research. Accounting, Organ. Soc. 27, 303–345 (2002)
43. Statistischer Bericht Berlin-Brandenburg, https://www.statistik-berlin-brandenburg.de/publikationen/Stat_Berichte/2012/SB_A01-05-00_2012h01_BE.pdf (Accessed: 14.11.14)
44. R Development Core Team: R: A Language and Environment for Statistical Computing, http://www.r-project.org/ (Accessed: 30.04.2014)
45. Verbeek, M.: A Guide to Modern Econometrics. Wiley, New York (2012)
46. Shapiro, S.S., Wilk, M.B.: An Analysis of Variance Test for Normality (Complete Samples). Biometrika. 52, 591–611 (1965)
47. Wooldridge, J.M.: Introductory Econometrics. A Modern Approach. South-Western, Mason (2013)
48. Breusch, T.S., Pagan, A.R.: A Simple Test for Heteroscedasticity and Random Coefficient Variation. Econometrica. 47, 1287–1294 (1979)
49. Koenker, R.: A Note on Studentizing a Test for Heteroscedasticity. J. Econom. 17, 107–112 (1981)
50. White, H.: A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity. Econometrica. 48, 817–838 (1980)