Monetary Incentive and Stock Opinions on Social Media

Hailiang Chen, Yu Jeffrey Hu, and Shan Huang

February 2019

Journal of Management Information Systems (JMIS), Forthcoming

Abstract

Social media not only is a new channel to obtain financial market information but also becomes the venue for investors to share and exchange investment ideas. We examine the performance consequences of providing monetary incentive to both existing and new amateur analysts on social media and its implications for online investor communities. We find that monetary incentive is effective in increasing the amount of content output and generating more interest from the community as well, but it leads to neither better nor worse stock recommendations. Additional analysis suggests that monetary incentive results in wider stock and industry coverage, a sign of increased content diversity. This study contributes to the understanding of the role of monetary incentive in stimulating the sharing of value-relevant information by investors in social media communities.

Keywords: monetary incentive, stock opinions, social media, wisdom of crowds, investment

1. Introduction

The advance of social media has made it possible for amateurs, who are not employed as professional security analysts and whose credentials are not endorsed by investment banks, to publish their stock opinions and recommendations to the public [2, 18]. Prior studies such as Chen et al. [9] have shown that stocks opinions shared on social media resemble analyst reports and can be useful in predicting future stock performance. While professional security analysts are incentivized by their high salaries, which are directly tied with their performance [16, 29], amateur analysts sharing their investment ideas openly online typically act independently and are not employed by social media platforms. Thus, most of the compensation mechanisms used by employers to induce effort from their workers such as wages, bonuses, and promotions [26] do not apply to these contributors on social media. To our knowledge, no prior study has examined whether monetary incentive can be a strong motivation for investors to share value-relevant information and produce better stock recommendations on social media. This study sets out to discover whether monetary incentive, in the form of ad revenue sharing by platform owners, can motivate platform contributors to generate a higher volume of content, and more importantly, improved quality of stock opinions and analyses.

Seeking Alpha (SA), one of the largest online crowd sourced equity research communities for investors, provides a unique setting to quantify the effect of monetary incentive on the quantity and quality of stock analysis and opinion articles. In January 2011, SA launched a premium partnership program that enables its contributors to earn $10 per 1,000 page views received by their "premium" articles, which are published exclusively on SA and not freely available anywhere else on the Internet.[1] After the implementation of this program, its contributors (both

[1] The premium partnership program was announced in an open letter from Seeking Alpha to contributors: https://seekingalpha.com/article/246803-an-open-letter-to-seeking-alpha-contributors. Over time, Seeking Alpha has made

existing and new) can continue to publish non-exclusive articles (we name them as "regular" articles) without monetary compensation on an article-by-article basis.

To investigate the effects of monetary incentive on the sharing of stock opinions on social media, we empirically test how the premium partnership program implemented by SA affects the content output of both existing and new contributors. Analyses for these two types of contributors could inform whether monetary incentive has a similar effect on different types of contributors, but more importantly, they also provide insights into whether and how monetary incentive may help increase the sustainability of online investor communities. Based on our rough estimation, 30% of SA contributors become inactive permanently within the first month after the publication of their first article; 70% of contributors quit contributing articles within a year. Thus, it is very important for online investor communities to evaluate whether offering monetary incentive helps retain existing contributors and at the same time attract new contributors in order to be sustainable.

For existing contributors, we employ the difference-in-differences (DID) approach to compare the articles published by regular and premium contributors in the period of two years both before and after the launch of the premium partnership program. We focus our analyses on existing contributors as of January 2011 who remained active and continued to publish articles after the launch of the program. The control group includes 241 contributors who published only regular articles in the study period. The treatment group includes 141 contributors who published only

several changes to the program. For instance, in June 2013, Seeking Alpha started to make a minimum payment of $150 for articles selected as a Small-Cap Insight; in July 2014, Seeking Alpha added an additional flat payment of $35 on top of the $10 per thousand page views. Details about these changes are available at https://seekingalpha.com/article/1475331-why-were-boosting-payments-to-high-value-contributors and https://seekingalpha.com/article/2343015-an-end-to-our-relationship-with-yahoo-a-new-era-for-equity-research?page=2. Our study period is 2009 to 2012, and thus our analysis is not affected by the changes introduced after 2012.

premium articles after the launch of the program.[2] The DID method is able to account for the inherent time-invariant differences between the two groups and estimate the effect of the treatment on the treated group.

The financial market context enables us to construct an objective quality measure, prediction accuracy, to assess the information content of articles published on SA by looking at whether the prediction in an article on a particular stock is consistent with its subsequent abnormal stock return [9]. This is one of the key differences between this study and prior empirical studies on the effects of monetary incentive on online word-of-mouth (WOM) or user-generated content (UGC), which mainly reflects consumers' personal opinions and lacks a good objective content quality measure.

Our analyses show that the number of articles on average increases by 68.8% after a contributor receives monetary payments from the platform. This is equivalent to 15 more articles published by a contributor in two years. However, the quality of the articles written after receiving monetary incentive does not significantly change in terms of accurately predicting the future price movements of the stocks covered. In other words, we find no evidence that monetary incentive can lead to either better or worse stock recommendations.

One issue in our context is that contributors voluntarily choose to participate in the program or not, which could potentially result in self-selection bias. To address this problem, we adopt a quasi-experiment design by performing propensity score matching [8, 13, 27] to identify a group of matched control and treatment contributors, who are observationally identical except the choice of participating in the premium partnership program. After matching, our results remain

[2] Only a small number of SA contributors publish both regular and premium articles. If we include these contributors in our analysis, each of them would be in the control group and treatment group simultaneously. This would make it difficult to interpret our results, so we exclude them from the analysis to err on the side of caution. However, our results do not change if all these contributors are assigned to either the treatment group or the control group.

qualitatively the same, although the effect size of monetary incentive on the number of published articles becomes smaller (i.e., an increase of 43.2%).

To shed more light on how monetary may influence the quality of content output, we examine two other measures related with article quality. As the compensation for each article is determined by the number of page views, one alternative quality measure is the interest level generated by each article. Although data about the number of page views is absent, we do observe all the comments received by each article. It is very likely that the number of comments is positively correlated with the number of page views per article. We find that premium contributors' articles on average receive 28.6% more comments than regular contributors' articles, implying that monetary incentive is effective in stimulating interest from the community. In addition, it is also interesting to investigate whether monetary incentive impacts the consistency level of a contributor's prediction accuracy, i.e., the variation of article quality. Our results, however, do not support this claim, and we find that monetary incentive does not lead to a higher variation in article quality for individual contributors.

We also conduct additional analysis on how monetary incentive changes the behavior of existing contributors in other content-related aspects. We first find that the average article length does not change after monetary incentive is provided. This result rules out the possibility that existing contributors write shorter articles in order to increase the quantity of articles. On a similar note, we find that the total number of words written by premium contributors increases significantly, which also suggests that existing contributors indeed exert more effort in the presence of monetary incentive. In addition, both the number of distinct stocks and the number of distinct industry sectors covered by an existing contributor increases significantly. Taken together, it is possible that existing contributors expand their scope of coverage by writing on more stocks from more industry sectors in order to publish more articles to receive monetary payoffs. In this regard, providing monetary incentive is effective in promoting content diversity for online investor communities.

We also investigate the behavior of new contributors after monetary incentive is introduced. We first show that the numbers of monthly new contributors are on average significantly larger after January 2011. This implies that the premium partnership program is successful in attracting new contributors to join the community. Among new contributors, more than two thirds of them are premium contributors; this proportion is much larger than that of existing contributors. To investigate how new contributors may behave differently, we also compare the content output of premium contributors with that of regular contributors. Our results for new contributors are largely consistent with the results for existing contributors. In short, we find that monetary incentive leads to more published articles for new premium contributors but does not have a significant effect on the article quality.

To our knowledge, this study is the first empirical research that investigates how offering monetary incentive to amateur analysts on social media affects their content output and quality of stock recommendations. Prior studies in the literature primarily focus on the value relevance of social media contents generated by investors and amateur analysts, but there is lack of research on how to better motivate them to share more value-relevant information. This study makes a first step to fill this gap. Our result that monetary incentive does not lead to better stock recommendations provides important implications for social media communities that facilitates the generation and spread of investment opinions by investors. Our findings suggest that monetary incentive is effective in promoting community engagement and to certain extent increasing content diversity but it does not necessarily encourage the sharing of more value-relevant information by investors on social media sites.

The remainder of this paper is organized as follows. In the next section, we provide a brief review of three research streams closely related with this study. We then introduce our dataset in Section 3. We present the empirical analyses and results for existing contributors and new contributors in Section 4 and Section 5, respectively. The last section concludes and discusses our limitations.

2. Literature Review

This study is first closely related to the literature on the role of social media in financial markets. Social media has significantly changed the way investors exchange information and share opinions about stocks in financial markets. Investors also increasingly rely more on investment advice shared on social media platforms, such as Twitter, Seeking Alpha, and Estimize. Earlier studies [1, 11, 34] in this literature examine how posts on Internet message boards affect stock returns and trading volumes. More recently, Bollen et al. [5] show that Twitter sentiments can predict the stock market index in the short run. Luo and Zhang [23] study the interrelationship between product reviews, website traffic, and firm value for information technology companies. Xu and Zhang [39] examine how information aggregation about public firms on Wikipedia may improve the information environment for investors in the financial market. With the increasing popularity of social media, researchers also examine whether social media contents reflect the wisdom of crowds [9, 17, 25] in stock market predictions. Chen et al. [9] use articles and comments posted on Seeking Alpha and find that stock opinions on social media can predict the future performance of individual stocks in the long run. Jame et al. [14] use crowdsourced earnings forecasts from Estimize and show that earnings forecasts posted by amateur analysts are incrementally useful in forecasting earnings and measuring the market's expectations. Xie et al. [38] investigate how network cohesion plays a role in affecting the prediction accuracy of social media analytics for financial markets. Mai et al. [24] study how social media impacts the value of Bitcoin using textual analysis and vector error correction models.

Our study differs from prior studies on the role of social media in financial markets by moving one step further and attempting to understand how monetary incentive provided by online community owners may affect the behavior of amateur analysts on social media and in particular whether it improves their performance, such as generating a greater amount of information that is of higher quality. Therefore, our study also builds upon the large economics literature that investigates the performance consequences of extrinsic rewards. Monetary incentive has been identified by prior studies as one of the most important motivations to promote effort and performance, along with intrinsic motivation and image motivation [3, 30, 33]. Although economic theories suggest that an increase in extrinsic rewards provided by activity can improve the effort and performance of participants, some studies [12, 15] show that contingent rewards can also be counterproductive, especially in the long run. One possible reason for this is that contributors could at the same time be motivated by other factors such as deriving inherent satisfaction through helping others (i.e., intrinsic motivation) and demonstrating their expertise and gaining social recognition (i.e., image motivation); when intrinsic and image motivations are crowded out by extrinsic incentives, contributors' outputs could thus decrease [3, 21].

Third, our study also adds to the growing information systems literature that examines how monetary incentive affects the quantity and quality of user-generated content (UGC). Previous studies in this area mainly focus on the context of online reviews and find mixed results. They show that monetary incentive increases the volume of customer reviews [7, 36], but it either is associated with reduced quality of customer reviews [19] or does not affect quality at all [36]. Liu and Feng [21] build a theoretical model to explain the contradictory results observed for the impact of monetary incentive on UGC. Wang et al. [37] also study the relationship between paid customer reviews and product sales. They find that when a retailer reduces monetary incentives for writing reviews, product sales decrease significantly.

Although prior studies in economics and information systems have documented how monetary rewards motivate individuals in various contexts, our study is among the first studies that examine the consequences of monetary incentive on social media content output for financial markets. The financial market context makes our study unique from the previous research on monetary incentive and UGC contribution. Customer reviews mainly reflect a customer's personal and subjective opinions based on the customer's experience with a product (i.e., there is no right or wrong), but stock opinions and recommendations will eventually be validated by the market (i.e., the quality of stock recommendations can be accurately assessed). Therefore, the results on the effects of monetary incentive on paid customer reviews may not generalize to the stock market context. Our study thus makes a significant contribution by empirically quantifying the effects of monetary incentive on stock opinions and recommendations on social media, which can be broadly considered as one type of user-generated content that aims to analyze and predict the future financial performance of public firms.

3. Data

Seeking Alpha is one of the biggest investment-related social media websites in the U.S. It had 7 million average monthly unique visitors in 2016 [28]. The website relies on a crowdsourced contributor network to publish analysis and opinion articles on a broad range of stocks including small- and mid-cap stocks. An editorial panel reviews all submitted articles and may provide feedback to improve clarification but not to interfere with the contributor's viewpoint. If deemed of adequate quality, these articles are then published on the SA website. In response to these articles, any interested user can write a commentary, sharing his or her own view, which may agree or disagree with the author's view on the stock in question.

We download all articles and comments posted on SA from 2009 to 2012. Each article is tagged with one or more stock tickers. Single-ticker articles focus solely on one stock, making it relatively easy to extract the author's opinion on that company. Multiple-ticker articles discuss more than one stock in the same article, rendering extraction of the author's various opinions for each of the tagged stocks difficult, if not impossible. We therefore focus our analysis primarily on single-ticker articles, for which we are able to clearly assess the prediction accuracy of contributors' opinions. Later we also conduct a robustness check on the effect of monetary incentive on content output that include both single-ticker and multi-ticker articles. The information we collect about each article includes the following items: article ID, title, main text, date of publication, author name, and stock ticker. We also extract all commentaries written in response to the single-ticker articles in our sample. The information we collect about each commentary includes the following items: article ID, comment ID, main text, date the comment is made, and author name. Whether the author of an article receives monetary compensation from SA or not is not public information. SA has provided us with proprietary data on whether an article is premium or regular (i.e., the author of a premium article receives payment from SA). For all the stocks covered in our sample, we collect financial-statement and financial market data from Compustat and the Center for Research in Security Prices (CRSP), respectively.

The sample period is set to be 2009 to 2012, which includes two years of data both before and after the event of introducing the premium partnership program (i.e., January 2011). For the sample of existing contributors, we restrict the sample to existing contributors as of January 2011 who remained active and continued to publish articles after the launch of the program. The control group includes 241 contributors who published only regular articles in the study period. The treatment group includes 141 contributors who published only premium articles after the launch of the program. For the sample of new contributors, we focus on the 1,045 new contributors who published their first articles in the two years after the event. There are 318 new regular contributors and 727 new premium contributors.

4. Effects of Monetary Incentive on Existing Contributors

We have two main dependent variables, Log($ArticleQuantity_{it}$) and $ArticleQuality_{it}$, which denote the quantity and quality of content output by contributor i in time period t, respectively. Log($ArticleQuantity_{it}$) is constructed as the natural logarithmic transformation of the number of single-ticker articles published by contributor i in time t. $ArticleQuality_{it}$ is measured as the percentage of single-ticker articles that correctly predict future stock returns. The prediction accuracy of each single-ticker article j, $Accuracy_j$, is calculated as follows. To quantify and study the views reflected in SA articles, we employ textual analysis. Specifically, we build on prior literature, suggesting that the frequency of negative words used in an article captures the tone of the report [9, 11, 20, 22, 31, 32]. We use the negative word list compiled by Loughran and McDonald [22] to characterize the views expressed in SA articles. An article is considered "bullish" if its fraction of negative words is below the median of its overall distribution across all contributors; an article is considered "bearish" if its fraction of negative words is above the median.[3] For each article in our sample, we compute the ensuing cumulative three-month abnormal return, which is the difference between the raw return and the return of a value-

[3] Different contributors may have different writing styles, and their level of use of negative words could also be different. We test the robustness of our results by applying contributor-specific medians when classifying articles into bullish or bearish ones, and our results remain the same (results are available upon request).

weighted portfolio of firms with similar size, book-to-market ratio, and past return [10].[4] There is a potential concern that SA articles may incite naïve investor reaction. That is, SA articles reflect false or spurious information yet still cause investors to trade in the direction of the underlying articles and move prices accordingly. To alleviate this concern, we skip the first two days after article publication when computing the cumulative abnormal return. $Accuracy_j$ equals to 1 if a bullish article is followed by positive abnormal returns OR if a bearish article is followed by negative abnormal returns. Otherwise, $Accuracy_j$ is set to be 0. $ArticleQuality_{it}$ is the sum of $Accuracy_j$ for all articles published by contributor i in time t divided by $ArticleQuantity_{it}$. Note that $ArticleQuality_{it}$ is expected to be 0.5 if a contributor without any valuable information simply does random guessing when predicting the sign of future abnormal returns.

4.1 Descriptive Analysis

To get an initial idea of the behavior of contributors on SA, we plot the average quantity and quality of the articles published by the existing contributors in the two groups. Figure 1 presents the average number of single-ticker articles per contributor in each month of the study period for both the control group and the treatment group. On average, existing contributors in the control group publish more articles than contributors in the treatment group in most months. In other words, on average relatively more productive contributors choose to not join the premium partnership program. Without monetary incentive, an average contributor in the control group tends to publish fewer articles over time in the entire study period. This is consistent with the general observation that many online contributors become less active or completely inactive after a certain time period [6, 35]. By contrast, an average contributor in the treatment group does not

[4] As a robustness check, we also construct the article quality measure using the cumulative one-month and six-month abnormal returns. Our results remain the same (See Table 5).

follow this decreasing trend but to some extent publishes relatively more articles after receiving monetary incentive.

[Figure 1]

Figure 2 presents the average article quality across existing contributors in two groups over time. The average quality is around half for both groups, although the variation of the treatment group is larger than that of the control group. Overall, Figures 1 and 2 together suggest that monetary incentive increases the content output produced by existing contributors but does not lead to an obvious improvement in their predictive accuracy. The curves prior to the event in these two figures also show a parallel pattern, which indicates that the parallel trend assumption holds in our context for the DID analysis.[5]

[Figure 2]

4.2 Monetary Incentive and Article Quantity/Quality

We now conduct the analysis under the regression framework. From Figure 1, we can learn that the behavior of regular contributors is potentially different from that of premium contributors. To explicitly control for the inherent differences between these two groups, we first employ the difference-in-differences (DID) approach and specify the following model as described in Equation (1). The idea of the DID method is that the effect of the treatment on the treated group can be captured by the change in the post-event differences in the outcome variables between the control and treatment groups compared with the pre-event differences. Following the suggestion by Bertrand et al. [4], we implement the standard two-period DID estimator instead of adopting a

[5] To formally test the parallel trend assumption under a regression framework, we estimate a model similar as Equation (1) on the data prior to the launch of the premium partnership program (year 2009 and 2010 only). We consider a placebo event that took place at the beginning of 2010 and define a variable *PlaceboAfter*, which is 1 if the year is 2010 and 0 if the year is 2009. The dependent variable is either the total number or the average accuracy of articles written by a contributor in a year. The coefficient estimates for the interaction term, *Treatment*×*PlaceboAfter*, are statistically insignificant at the 10% level in both regressions, which suggests that the trends of the dependent variables follow a parallel trend for the two groups prior to the launch of the premium partnership program.

contributor/month panel (i.e., there are multiple monthly observations in both pre-event and post-event periods) to avoid the downward biased estimation of the standard errors for model coefficients, resulting from serial correlations. More importantly, many contributors may not write any article in certain months, which makes it impossible to assess one of the outcome variables (article quality) on a monthly basis. Thus, it is reasonable to aggregate the monthly observations over a longer period to avoid zero observations for article quantity and missing values for article quality.

$$Y_{it} = \beta_0 + \beta_1 Treatment_i \times After_t + \beta_2 After_t + \beta_3 Log(ActivePeriod_{it}) + \beta_4 Log(DaysOnSA_{it}) + f_i + \varepsilon_{it} \tag{1}$$

The dependent variable is either Log($ArticleQuantity_{it}$) or $ArticleQuality_{it}$. $Treatment_i$ is 1 for existing contributors in the treatment group and 0 for existing contributors in the control group. $After_t$ is 0 for t=1 (before the launch of the premium program) and 1 for t=2 (after the launch of the premium program). $Treatment_i \times After_t$ is the main variable of interest. $ActivePeriod_{it}$ is the number of days contributor *i* is active in time period *t*. Note that our sample of contributors include only existing contributors at the time of launching the premium partnership program. In the period before the event, if a contributor published an article before the study period, then $ActivePeriod_{it}$ is the length of the pre-event period. If a contributor started to publish the first article in the middle of the pre-event period, $ActivePeriod_{it}$ is the number of days between the publication date of the first article and the end of the pre-event period. In the period after the event, $ActivePeriod_{it}$ is the length of the post-event period if the contributor continued to publish articles after the end of the study period. If a contributor stopped to publish articles in the middle of the post-event period, $ActivePeriod_{it}$ is the number of days between the beginning of the post-event period and the publication date of the last article by the contributor. $DaysOnSA_{it}$ is the number of days between the date of contributor i's first article on SA and the date when contributor i published the last article in time period t. This variable controls for the effect of a contributor's tenure on SA. f_i is the individual fixed effects, which absorbs the time-invariant variable $Treatment_i$. $\mathcal{E}_{it}$ is the error term.

One issue to be addressed is the self-selection problem that could lead to potential endogeneity. Different from a randomized experimental design, premium contributors in our context self-selected to participate in the premium partnership program instead of being randomly assigned to it. We adopt a quasi-experimental design and perform propensity score matching [8, 13, 27] to identify a group of matched regular contributors that are observationally identical to the treated premium contributors except the treatment condition. For this purpose, we use the following 14 observable variables constructed as of the day prior to the event day to characterize an existing contributor's choice of participating in the premium partnership program.

$DaysOnSA_i$ is the number of days from the publication date of a contributor's first SA article to the event date. This variable measures an existing contributor's tenure on SA, or how long an amateur analyst has been active and sharing stock investment ideas with others on social media. $Follower_i$ is the number of followers a contributor has on SA prior to the event and is a proxy of how popular a contributor is in the SA community. $CommentsPerArticle_i$ is the average number of comments received by a contributor's articles published prior to the event. To enjoy the benefits of the premium partnership program, a contributor's articles need to have a significant amount of page views. The number of page views received by each article published before the event is not available, but the number of comments on each article is likely to be positively correlated with the number of page views. This variable thus quantifies the potential monetary compensation a contributor can receive on average by writing each additional premium article. $NegFraction_i$ and $StdevNegFraction_i$ are the average and standard deviation of the fraction of negative words across all articles published prior to the event by contributor *i*, respectively. These two variables related with article sentiment capture a contributor's writing style and also potentially the amount of private information about the stocks they cover. $NegFraction_i$ reflects on average how negative a contributor's article is compared with all other articles. $StdevNegFraction_i$ shows the variation in the sentiments of different articles written by the same contributor. When this variable is large, it implies that the contributor tends to cover stocks with different levels of performance. $Company_i$ is an indicator variable denoting whether a contributor is affiliated with a company as disclosed on SA. $Blog_i$ is an indicator variable denoting whether a contributor maintains a blog site. $Company_i$ and $Blog_i$ are intended to capture a contributor's activities outside SA given that one condition for receiving payment from SA is its exclusive right of publishing premium articles. The remaining seven variables are dummy variables generated from the occupation of a contributor: $Academic_i$ (students and professors), $Advisor_i$, $Analyst_i$, $Executive_i$, $FundManager_i$, $InvestorTrader_i$, and $Journalist_i$ (including journalists, newsletter authors, and professional bloggers). The omitted category of other or unknown occupations serves as the baseline case. Investors of different occupation are likely to have different motivations for sharing their investment ideas on social media. Thus, we also expect that there would be considerable heterogeneity across different occupations in their decision of joining the premium partnership program.

[Table 1]

Panel A and Panel B of Table 1 present the summary statistics of the dependent/independent variables for regression analyses and contributor characteristics used for propensity score matching described above. The number of observations is 764 for dependent and independent variables constructed for both pre- and post-event periods. The mean of *ArticleQuantity* (without log transformation) is 22.13, which is equivalent to about one article per month for each contributor. The variation in the number of published articles across contributors is large, as the standard deviation is 107.24. The median number of published articles in two years is only 4, but the maximum number reaches 2,145. The mean of *ArticleQuantity* is 0.51, which is slightly larger than 0.5, the expected average with no valuable information. This indicates that on average the opinions revealed on SA may contain value-relevant information about the long-term performance of stocks. The standard deviation is 0.31, suggesting that there is a significant variation in the predictive accuracy of articles written by different contributors. The mean of *ActivePeriod* is 512.29 days, which is roughly 17 months. The median is 639.50 days (21 months) and larger than the mean, suggesting that most contributors are active almost over the entire study period. However, there are a few contributors that abandoned the site shortly after the launch of the program, as the minimum of *ActivePeriod* is only 1 day.

In Panel B, the number of observations is 382 for the variables related with contributor characteristics as of the end of the pre-event period. The mean and median of an existing contributor's tenure on SA (*DaysOnSA*) are 686.53 and 478.91 days, respectively. The maximum can reach 2,234 days (more than six years). The size of follower networks on SA is not large compared to other popular social media sites such as Twitter, as the mean of *Follower* is 3,721 and the median is only 132. The mean and median of *CommentsPerArticle* are 8.58 and 6.50, respectively. The minimum is 0 but the maximum reaches 61.85. The average fraction of negative words is 1.4% and its standard deviation is 0.5%. As to the occupations of contributors, the majority of contributors are investors or traders (24%), journalists (19%), and analysts (13%).

[Table 2]

We report the results of the Probit regressions conducted before and after matching in Table 2. The dependent variable is an existing contributor's choice of whether to join the premium partnership program. In Column (1), the regression is performed on the full sample of 382 existing contributors. We find that contributors with longer tenure on SA, less followers, more comments per article, and smaller standard deviation of article sentiment are more likely to join the program and receive monetary payments. Contributors from the occupations of analysts and investors/traders are also more likely to join the treatment group. After the matching based on Nearest Neighbor with one neighbor and with common support, we identify 212 matched regular and premium contributors that achieve balance in all observables. Column (2) of Table 2 presents the Probit regression result on this matched sample. All the co-variates are statistically insignificant at the 10% level, confirming that the matched control and treatment groups are observationally identical except for the treatment condition.

[Table 3]

Table 3 presents the estimation results of Equation (1) for the effects of monetary incentive on article quantity and quality. The dependent variable in Columns (1) and (2) is Log($ArticleQuantity_{it}$) and the dependent variable in Columns (3) and (4) is $ArticleQuality_{it}$. Columns (1) and (3) report the results of the DID analysis based on the full sample of 382 existing contributors, while Columns (2) and (4) report the results of the DID analysis based on the matched sample of 212 existing contributors.

The coefficient estimate on $Treatment_i \times After_t$ in Column (1) is 0.688 and statistically significant at the 1% level, suggesting that the number of articles published by a contributor on average increased by 68.8% after monetary payments are provided for writing articles. The result for the matched sample also yields a positive effect of monetary incentive, but the effect size decreases to 0.432 in Column (2), indicating that the number of articles published by a contributor on average increased by 43.2% after being motivated by monetary payments. However, we observe very different results when the dependent variable is the article quality. The coefficient estimate on $Treatment_i \times After_t$ in Column (3) is positive but statistically insignificant at the 10% level. In Column (4), we also observe an insignificant coefficient estimate on $Treatment_i \times After_t$ for the matched sample. These results suggest that monetary incentive is effective in driving up the quantity of content outputs by existing contributors, but it does not lead to the contribution of articles of higher quality.

The results for control variables in Table 3 are consistent with expectations. The coefficient estimate on *After* is negative and statistically significant at the 1% level in both Columns (1) and (2). This is consistent with the decreasing pattern revealed in Figure 1, which again indicates that the productivity of a contributor decreases over time. The coefficient estimate on *After* is positive but statistically insignificant at the 10% level in both Columns (3) and (4). This suggests that the quality of articles written in the post-event period is on average similar as that of articles written in the pre-event period. The coefficient estimate on *ActivePeriod* is positive and statistically significant at the 1% level in both Columns (1) and (2), because the longer a contributor is active on the SA site in each time period, the more articles the contributor writes. The coefficient estimate on *DaysOnSA* is negative and statistically significant at the 5% level in Column (4), suggesting that article quality tends to be lower for contributors with a longer tenure on SA.

To shed more light on the relationship between monetary incentive and article quality, we conduct additional analyses on two other measures associated with article quality in Columns (5) to (8) of Table 3. *Log(CommentsPerArticle)* is the natural logarithmic transformation of the average number of comments received by an article, which reveals the interest level of an article and can thus serve as an alternative measure for article quality. *ArticleQualityVariation* is the standard deviation of *Accuracy* across all articles for a contributor and thus measures the variation of a contributor's article quality. The coefficient estimate on $Treatment_i \times After_t$ in Column (6) is 0.286 and statistically significant at the 5% level, suggesting that premium contributors' articles on average receive 28.6% more comments than regular contributors' articles. The coefficient estimate on $Treatment_i \times After_t$ in Column (8) is, however, statistically insignificant at the 10% level, suggesting that monetary incentive does not have a significant effect on the variation of article quality in terms of prediction accuracy. In sum, the results from these additional analyses further imply that premium articles can generate a higher level of interest albeit are not more accurate in predicting future returns of the stocks discussed.

4.3 Monetary Incentive and Article Content

To further investigate how monetary incentive changes the behavior of contributors, we define three other content-related variables and conduct a similar set of analyses as in Table 3 but replace the dependent variables with the following new variables. $ArticleLength_{it}$ is the average number of words in articles written by contributor i in time period t. $TickerCoverage_{it}$ is the number of unique stocks covered by articles written by contributor i in time period t. $SectorCoverage_{it}$ is the number of unique industry sectors covered by articles written by contributor i in time period t. Article length measured in number of words can be another measure of a contributor's effort in sharing their investment views and opinions in addition to the simple count of article numbers. It is possible that contributors may write shorter but more articles in response to monetary incentive assuming that it takes the same amount of time and effort to write the same number of words. As to the ticker and sector coverage, a key constraint each contributor faces is that it takes a significant amount of time and effort for any contributor to get familiar with the fundamentals of a public company and gather the relevant information about an industry sector. In light of this, a contributor can potentially adopt two approaches: one is to cover a few tickers/sectors but write as many articles as possible on each target company or industry sector, and the other is to cover more tickers/sectors but write only a few articles on each. Thus, by studying how monetary incentive affects the ticker/sector coverage, we can distinguish between the two approaches adopted by contributors for the same purpose of writing more articles.

The summary statistics of these three variables are provided in Panel C of Table 1. The average length of SA articles is 818.38 words. The standard deviation is 472.79, implying that there is a significant level of variation in the length of different articles. The mean and median number of stocks covered by a contributor in each time period is 10.09 and 3, respectively. The maximum number of stocks covered can reach 648. The mean and median number of industry sectors covered by a contributor is 2.64 and 2, respectively, while the maximum reaches 9. These statistics suggest that on average a contributor specializes in a few industry sectors and covers a small range of firms.

[Table 4]

Table 4 reports the results for how monetary incentive affects the length of published articles and the coverage of stocks and industry sectors. Similarly as in Table 3, we also report the results from the difference-in-differences model for both the full and matched samples. The results are largely consistent between these two samples. The effects based on the full sample, when statistically significant, are slightly larger in magnitude than that of the matched sample. Specifically, the length of articles written by existing contributors does not change significantly after monetary incentive is offered as shown in both Columns (1) and (2). While the result in Table 3 suggests that a contributor publishes more articles with the presence of monetary incentive, the result in Table 4 further implies that contributors do not intentionally cut down the length of articles in order to publish more articles. Based on these results, we can infer that existing contributors indeed put more effort into the task of writing more articles after monetary incentive is introduced. In addition, results in Columns (3) to (6) suggest that a contributor covers more stocks and industry sectors after receiving monetary incentive. In summary, these findings together suggest that contributors may push themselves to study more firms and expand to new industry sectors in order to write and publish more articles.

Although not our main interest, the coefficient estimates on *After* in Table 4 reveal some interesting patterns about the differences between the post-event and pre-event periods. On average, the length of articles does not significantly differ between the pre- and post-event periods (Column 2). Contributors cover significantly less stocks and less industry sectors in the post-event period than in the pre-event period (Columns 3 to 6). These results can be attributed to certain community-wide changes over time that impact both regular and premium contributors. In addition, article length is not associated with a contributor's duration of active period, but the longer a contributor stays active in each time period, the more tickers he or she covers. A contributor's tenure on SA is also positively associated with both ticker and industry sector coverage.

4.4 Robustness Checks

We conduct three robustness checks to further validate our results. First, we adopt the total word count of all articles published by each contributor to measure the quantity of content output instead of the article count. This measure considers both article count and average word count per article at the same time. We find that monetary incentive significantly increases the total number of words written by a contributor (see Columns 1 and 2 of Table 5). Second, we try to use a different time window when evaluating the future abnormal return and constructing the article quality measure. Our results remain robust (i.e., monetary incentive leads to neither better nor worse stock recommendations) when either one month (Columns 3 and 4 of Table 5) or six months (Columns 5 and 6 of Table 5) is selected. Third, our analyses thus far focus on single-ticker articles only, because it is difficult to assess the sentiment associated with each of the stock tickers discussed in a multi-ticker article. However, it is possible that monetary incentive may induce contributors to write more single-ticker articles and less multi-ticker articles, but the total number of articles may not change. To address this concern, in Table 6 we conduct a robustness check by incorporating multi-ticker articles into our analysis and examining the quantity, article length, ticker coverage and industry sector coverage of the overall content output (i.e., combining single-ticker and multi-ticker articles). The results reported in Table 6 are quite consistent with those for single-ticker articles only. We continue to find that monetary incentive increases the overall quantity of articles (including both single-ticker and multi-ticker ones) written by a contributor (Column 2 of Table 6); monetary incentive does not affect the average article length (Column 4); and monetary incentive increases both stock coverage and industry sector coverage (Columns 6 and 8).

[Table 5]

[Table 6]

5. Effects of Monetary Incentive on New Contributors

While studying the behavior change of existing contributors on SA allows us to make a before-and-after comparison, it does not inform us how the presence of monetary incentive may attract new contributors and affect their content contribution behaviors. To examine whether the introduction of the premium partnership program attracted many new contributors, we plot the numbers of new contributors in each month over the period of 2009 to 2012 in Figure 3 (solid curve). In each month, we count the number of contributors whose first articles were published in that month. Each contributor is counted only once over time. We can see that the numbers of monthly new contributors increased significantly in the post-event period compared with that of the pre-event period. For comparison, we also plot the numbers of new regular contributors after the event in Figure 3 (dotted curve). The numbers of new regular contributors after the event appear to be roughly similar as the numbers of new contributors prior to the event, when there were only regular contributors. These patterns suggest that monetary incentive is quite effective in attracting new contributors to join the community.

To further investigate whether our results from the analyses on existing contributors generalize to new contributors, we conduct an additional set of analyses focusing on the sample of 1,045 new SA contributors that joined the community in the two years after the event. Among them, 318 contributors belong to the control group and 727 contributors belong to the treatment group. Since only data for the post-event period is available for these new contributors, the DID analysis cannot be applied on this sample. Many of the contributor characteristics described in Panel B of Table 1 cannot be constructed unless we use ex post information. For this reason, we also refrain from applying the propensity score matching procedure to obtain a matched sample of regular and premium contributors. Therefore, we run simple OLS regressions on the sample of all 1,045 new contributors and assess how the participation in the premium partnership program affects the content output of new contributors. We acknowledge that this analysis may suffer from the self-selection bias concern and thus should be interpreted with caution.

[Table 7]

Table 7 presents the summary statistics of the variables for the new regular and premium contributors in two panels separately. For new contributors, the number of days on SA is the same as the active period, so we no longer control for *DaysOnSA* in the analysis. New premium contributors on average write more articles, have a higher variation in article quality, and cover more stocks and more industry sectors than new regular contributors. However, the article quality of new premium contributors is slightly lower than that of new regular contributors, although the difference is insignificant according to a simple t-test.

[Table 8]

Table 8 presents the results of the OLS regressions. The results for new contributors are largely consistent with the results obtained for existing contributors. Specifically, we find that new premium contributors write 32.6% more articles than new regular contributors, but there is no significant difference in article quality between two groups of new contributors. New premium contributors' articles receive 35.7% more comments than new regular contributors' articles. The variation in article quality is also 9.3% higher for new premium contributors. Article length does not differ between the two groups, but new premium contributors cover 27.6% more stocks and 16.6% more industry sectors than new regular contributors.

6. Discussion and Conclusion

Social media is playing an increasingly important role in financial markets. It not only is the new and additional channel to obtain market information besides traditional media outlets (e.g., newspapers and TV) but also becomes the venue for investors to share and exchange investment ideas. This study contributes to the literature on the role of social media in financial markets by providing empirical evidence on how monetary incentive offered by an online investor community affects the quantity and quality of content outputs by amateur analysts on social media. We find that monetary incentive increases content contribution and diversity but does not lead to improved quality of stock recommendations.

Our results provide important managerial implications for various social media sites. As monetary incentive mechanisms such as ad-revenue-sharing programs are becoming more popular, social media community owners should carefully evaluate the implications of these programs and align them with their goals. In particular, given our result that monetary incentive does not lead to improved prediction accuracy but nevertheless generates more interests as reflected in article comments, we cannot rule out the possibility that some online contributors who accept payments from the social media community may not have valuable private information to share or even act with the purpose of attracting attention only.

There are a few limitations with our study. First, to infer causality in an observational study, we employ the quasi-experimental design and propensity score matching to address the potential selection bias and endogeneity issues. Future research can further validate our results using other methodologies such as randomized experiments. Second, the monetary compensation offered in our context is relatively small, especially compared with the high salary received by professional security analysts. It is possible that the effect of monetary incentive on amateur analysts may depend on the amount of monetary gain.

References

1. Antweiler, W., and Frank, M. Z. Is all that talk just noise? The information content of internet stock message boards. *Journal of Finance*, 59, 3 (2004), 1259–1294.
2. Aral, S., Dellarocas, C., and Godes, D. Introduction to the special issue—Social media and business transformation: A framework for research. *Information Systems Research*, 24, 1, (2013), 3-13.
3. Benabou, R., and Tirole, J. Intrinsic and extrinsic motivation. *The Review of Economic Studies*, 70, 3, (2003), 489–520.
4. Bertrand, M., Duflo, E., and Mullainathan, S. How much should we trust difference-in-difference estimates. *The Quarterly Journal of Economics*, 119, 1, (2004), 249-275.
5. Bollen, J., Mao, H., and Zeng, X. Twitter mood predicts the stock market. *Journal of Computational Science*, 2, 1, (2011), 1-8.
6. Brandtzaeg, P. B., and Heim, J. User loyalty and online communities: Why members of online communities are not faithful. *Proceedings of the 2nd international conference on INtelligent TEchnologies for interactive enterTAINment (INTETAIN)*, Cancun, Mexico, 2008.
7. Burtch, G., Hong, Y., Bapna, R., and Griskevicius, V. Stimulating online reviews by combining financial incentives and social norms. *Management Science*, forthcoming, (2017).
8. Caliendo, M., Kopeining, S. Some practical guidance for the implementation of Propensity Score Matching. *Journal of Economic Surveys*, 22, 1, (2008), 31-72.
9. Chen, H., De, P., Hu, Y. J., and Hwang, B.-H. Wisdom of Crowds: The value of stock opinions transmitted through social media. *Review of Financial Studies*, 27, 5, (2014), 1367–1403.
10. Daniel, K., Grinblatt, M., Titman, S., and Wermers, R. Measuring mutual fund performance with characteristic-based benchmarks. *Journal of Finance*, 52, 3, (1997), 1035-1058.

11. Das, S. R., and Chen, M. Y. Yahoo! for Amazon: Sentiment extraction from small talk on the web. *Management Science*, 53, 9, (2007), 1375–1388.

12. Deci, E., Koestner, R., and Ryan, R. A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. *Psychological Bulletin*, 125, 6, (1999), 627–668.

13. Dehejia, R.H. and Wahba, S. Propensity Score-Matching methods for nonexperimental causal studies. *Review of Economics and Statistics*, 84, 1, (2002), 151-161.

14. Jame, R., Johnston, R., Markov, S., and Wolfe, M. C. The value of crowdsourced earnings forecasts. *Journal of Accounting Research*, 54, 4, (2016), 1077-1110.

15. Gneezy, U., and Rustichini, A. Pay enough or don't pay at all. *The Quarterly Journal of Economics*, 115, 3, (2000), 791–810.

16. Groysberg, B., Healy, P. M., and Maber, D. A. What drives sell-side analyst compensation at high-status investment banks? *Journal of Accounting Research*, 49, 4, (2011), 969-1000.

17. Hill, S., and Ready-Campbell, N. Expert stock picker: The Wisdom of (experts in) Crowds. *International Journal of Electronic Commerce*, 15, 3, (2011), 73-102.

18. Kane, G. C., Alavi, M., Labianca, G. J., and Porgatti, S. P. What's different about social media networks? A framework and research agenda. *MIS Quarterly*, 38, 1, (2014), 275-304.

19. Khern-am-nuai, W., Kannan, K., and Ghasemkhani, H. Extrinsic versus intrinsic rewards for contributing reviews in an online platform. *Information Systems Research*, forthcoming, (2017).

20. Li, F. Annual report readability, current earnings, and earnings persistence. *Journal of Accounting and Economics*, 45, 2-3, (2008), 221–47.

21. Liu, Y., and Feng, J. Does money talk? The impact of monetary incentives on User-Generated Content contributions. Working Paper, 2016.

22. Loughran, T., and McDonald, B. When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *Journal of Finance*, 66, 1, (2011), 35–65.

23. Luo, X., and Zhang, J. How do consumer buzz and traffic in social media marketing predict the value of the firm? *Journal of Management Information Systems*, 30, 2, (2013), 213-238.

24. Mai, F., Shan, Z., Bai, Q., Wang, X., and Chiang, R. H. L. How does social media impact Bitcoin value? A test of the silent majority hypothesis. *Journal of Management Information Systems*, 35, 1, (2018), 19-52.

25. Nofer, M., and Hinz, O. Are crowds on the internet wiser than experts? The case of a stock prediction community. *Journal of Business Economics*, 84, 3, (2014), 303-338.

26. Prendergast, C. The provision of incentives in firms. *Journal of Economic Literature*, 37, 1, (1999), 7-63.

27. Rosenbaum, P. R., and Rubin, D. B. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70, 1, (1983), 41–55.

28. Seeking Alpha. *About Seeking Alpha*. Available from http://seekingalpha.com/page/about_us. 2017.

29. Stickel, S. E. Reputation and performance among security analysts. *Journal of Finance*, 47, 5, (1992), 1811–1836.

30. Tang, Q., Gu, B., Whinston, A. B. Content contribution for revenue sharing and reputation in social media: A dynamic structural model. *Journal of Management Information Systems*, 29, 2, (2012), 41-76.

31. Tetlock, P. C. Giving content to investor sentiment: The role of media in the stock market. *Journal of Finance*, 62, 3, (2007), 1139–1168.

32. Tetlock, P. C., Saar-Tsechansky, M., and Macskassy, S. More than words: Quantifying language to measure firms' fundamentals. *Journal of Finance*, 63, 3, (2008), 1437–1467.

33. Toubia, O., and Stephen, A. T. Intrinsic vs. image-related utility in social media: Why do people contribute content to Twitter? *Marketing Science*, 32, 3, (2013), 368–392.

34. Tumarkin, R., and Whitelaw, R. F. News or noise? Internet postings and stock prices. *Financial Analysts Journal*, 57, 3, (2001), 41–51.

35. Velasquez, A., Wash, R., Lampe, C., and Bjornrud, T. Latent Users in an Online User-Generated Content Community. *Computer Supported Cooperative Work,* 23, 1, (2014), 21-50.

36. Wang, J., Ghose, A., and Ipeirotis, P. Bonus, disclosure, and choice: What motivates the creation of high-quality paid reviews? *International Conference on Information Systems (ICIS)*, Orlando, Florida, USA. 2012.

37. Wang, S., Pavlou, P. A., and Gong, J. On monetary incentives, online product reviews, and sales. *International Conference on Information Systems (ICIS)*, Dublin, Ireland. 2016.

38. Xie, P., Chen, H., and Hu, Y. J. Network structure and predictive power of social media for the Bitcoin market. Working Paper. 2017.

39. Xu, S. X., and Zhang, X. Impact of Wikipedia on market information environment: Evidence on management disclosure and investor reaction. *MIS Quarterly*, 37, 4, (2013), 1043-1068.

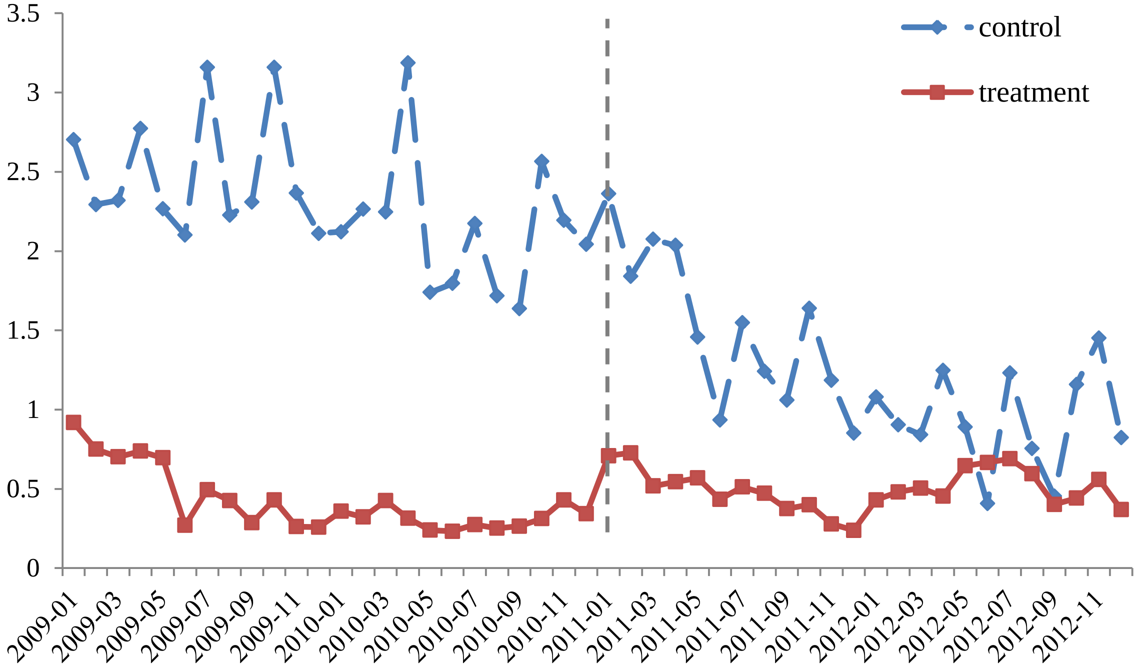

Figure 1. Average number of articles per contributor in each month

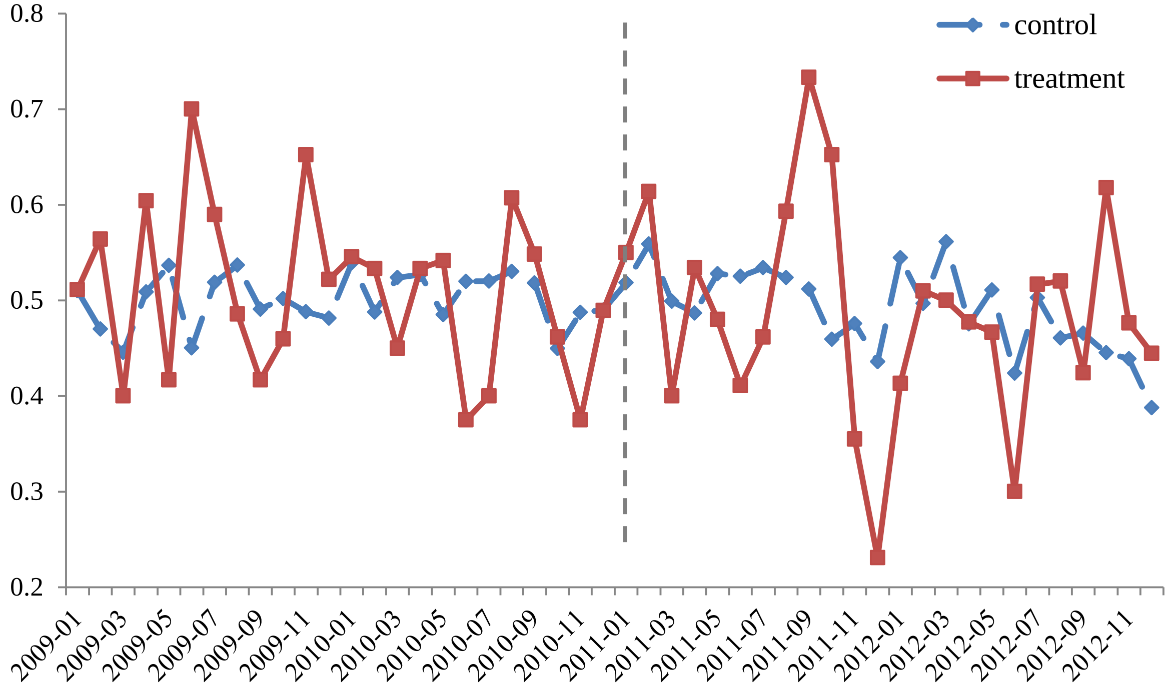

Figure 2. Average article quality in each month

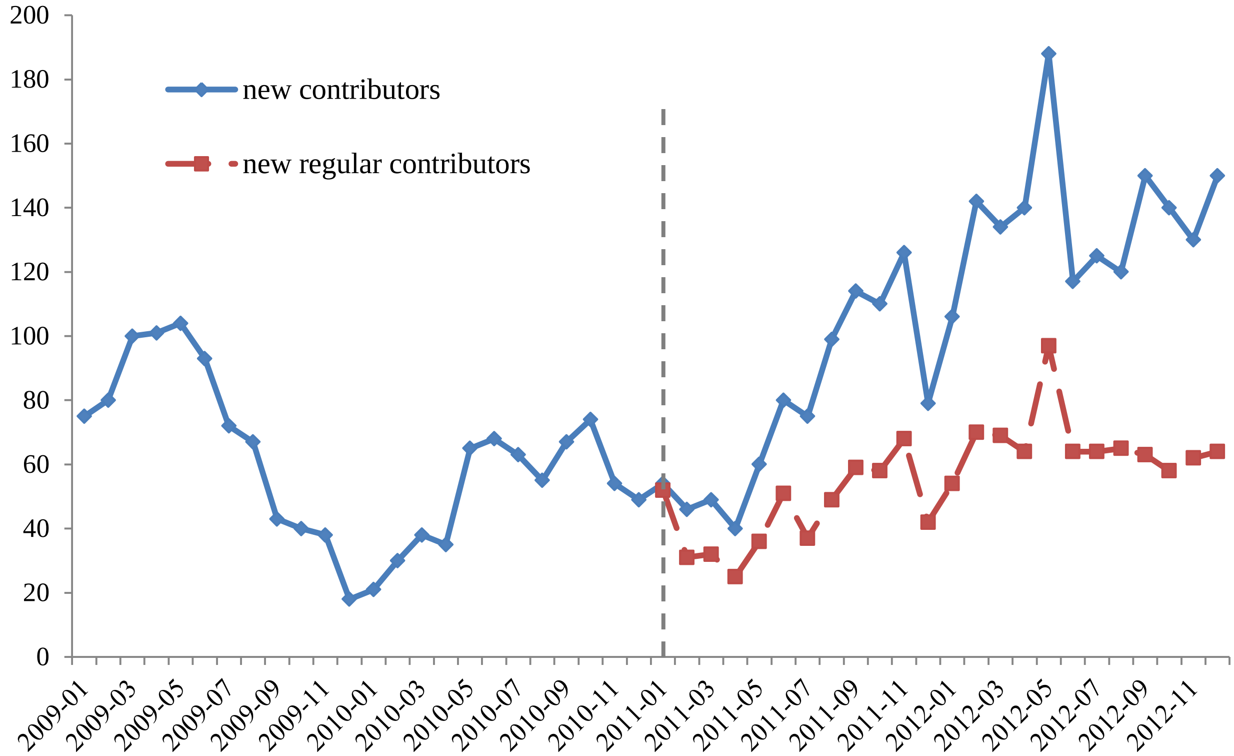

Figure 3. Number of new contributors in each month

Table 1. Summary Statistics: Existing Contributors

Variable	#Obs	Mean	Std. Dev.	Min	Median	Max
Panel A: Key Dependent and Independent Variables for Regressions						
ArticleQuantity	764	22.13	107.24	1	4	2,145
ArticleQuality	764	0.51	0.31	0	0.50	1
CommentsPerArticle	764	10.02	13.23	0	6	137.29
ArticleQualityVariation	764	0.31	0.23	0	0.47	0.5
ActivePeriod	764	512.29	240.97	1	639.50	711
DaysOnSA	764	871.44	556.33	0	808	2,817
Panel B: Variables for Propensity Score Matching						
DaysOnSA	382	686.53	478.91	0	624	2,234
Followers	382	3,720.60	12,726.35	0	131.50	84,579
CommentsPerArticle	382	8.58	7.68	0	6.50	61.85
NegFraction	382	0.014	0.008	0.001	0.012	0.050
StdevNegFraction	382	0.005	0.004	0	0.005	0.019
Company	382	0.49	0.50	0	0	1
Blog	382	0.41	0.49	0	0	1
Academic	382	0.03	0.17	0	0	1
Advisor	382	0.08	0.27	0	0	1
Analyst	382	0.13	0.34	0	0	1
Executive	382	0.05	0.21	0	0	1
FundManager	382	0.07	0.26	0	0	1
InvestorTrader	382	0.24	0.43	0	0	1
Journalist	382	0.19	0.39	0	0	1
Panel C: Content Variables						
ArticleLength	764	818.38	472.79	43	714.86	6,004
TickerCoverage	764	10.09	36.57	1	3	648
SectorCoverage	764	2.64	2.04	1	2	9

Table 2. Probit Regressions for Propensity Score Matching

	Before Matching	*After Matching*
	(1)	(2)
Constant	-1.413***	-0.246
	(0.549)	(0.620)
Log(DaysOnSA)	0.133*	0.055
	(0.073)	(0.082)
Log(Followers)	-0.196***	-0.056
	(0.040)	(0.049)
Log(CommentsPerArticle)	0.408***	0.002
	(0.115)	(0.139)
NegFraction	9.407	-10.601
	(9.593)	(11.878)
StdevNegFraction	-69.082***	21.453
	(21.809)	(26.657)
Company	0.275	0.327
	(0.170)	(0.203)
Blog	-0.214	-0.104
	(0.160)	(0.196)
Academic	0.184	0.040
	(0.467)	(0.548)
Advisor	0.382	0.132
	(0.335)	(0.411)
Analyst	0.527*	0.155
	(0.283)	(0.353)
Executive	-0.143	0.315
	(0.414)	(0.556)
FundManager	0.510	0.471
	(0.323)	(0.402)
InvestorTrader	1.197***	0.159
	(0.248)	(0.312)
Journalist	0.394	0.032
	(0.262)	(0.329)
Log Likelihoold	-196.8	-143.5
Observations	382	212

Note: *** p<0.01, ** p<0.05, * p<0.

Table 3. Impact of Monetary Incentive on Article Quantity and Quality

	Log(ArticleQuantity)		*ArticleQuality*		*Log(CommentsPerArticle)*		*ArticleQualityVariation*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment × After	0.688***	0.432***	0.030	0.031	0.210*	0.286**	0.066**	0.012
	(0.096)	(0.123)	(0.048)	(0.068)	(0.113)	(0.139)	(0.029)	(0.038)
After	-0.709***	-0.432***	0.020	0.079	0.154	0.173	-0.086***	-0.023
	(0.090)	(0.119)	(0.044)	(0.070)	(0.110)	(0.158)	(0.025)	(0.039)
Log(ActivePeriod)	0.209***	0.233**	0.033	0.075	0.011	-0.002	0.045**	0.053**
	(0.066)	(0.100)	(0.033)	(0.046)	(0.077)	(0.092)	(0.018)	(0.025)
Log(DaysOnSA)	0.368***	0.234	-0.065	-0.167**	0.071	0.073	0.069**	0.045
	(0.127)	(0.157)	(0.059)	(0.074)	(0.146)	(0.178)	(0.033)	(0.041)
Specification	DID+FE	DID+FE	DID+FE	DID+FE	DID+FE	DID+FE	DID+FE	DID+FE
Authors	382	212	382	212	382	212	382	212
Observations	764	424	764	424	764	424	764	424
R^2	0.352	0.270	0.003	0.024	0.083	0.139	0.157	0.136

Note: *** p<0.01, ** p<0.05, * p<0.1.

Table 4. Impact of Monetary Incentive on Article Length and Stock / Sector Coverage

	Log(ArticleLength)		*Log(TickerCoverage)*		*Log(SectorCoverage)*	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment × After	-0.063	0.009	0.475***	0.292***	0.246***	0.125**
	(0.050)	(0.066)	(0.080)	(0.104)	(0.049)	(0.064)
After	0.061*	-0.036	-0.545***	-0.355***	-0.249***	-0.168***
	(0.036)	(0.060)	(0.071)	(0.100)	(0.043)	(0.063)
Log(ActivePeriod)	0.005	-0.024	0.109**	0.118*	0.039	0.023
	(0.029)	(0.043)	(0.044)	(0.070)	(0.028)	(0.045)
Log(DaysOnSA)	0.052	0.105	0.374***	0.295**	0.219***	0.233***
	(0.048)	(0.065)	(0.097)	(0.118)	(0.056)	(0.074)
Specification	DID+FE	DID+FE	DID+FE	DID+FE	DID+FE	DID+FE
Authors	382	212	382	212	382	212
Observations	764	424	764	424	764	424
R^2	0.038	0.023	0.306	0.238	0.233	0.214

Note: *** p<0.01, ** p<0.05, * p<0.1.

Table 5. Robustness Checks

	Log(TotalWordCount)		*ArticleQuality-1-month*		*ArticleQuality-6-month*	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment × After	0.754***	0.507***	0.027	0.031	-0.067	-0.082
	(0.130)	(0.170)	(0.051)	(0.071)	(0.048)	(0.066)
After	-0.805***	-0.567***	0.039	0.099	-0.043	-0.006
	(0.117)	(0.168)	(0.046)	(0.076)	(0.041)	(0.062)
Log(ActivePeriod)	0.255***	0.269**	0.003	0.052	0.011	0.012
	(0.079)	(0.126)	(0.035)	(0.054)	(0.035)	(0.049)
Log(DaysOnSA)	0.532***	0.413**	-0.090	-0.186**	0.039	0.009
	(0.163)	(0.208)	(0.060)	(0.085)	(0.056)	(0.072)
Specification	DID+FE	DID+FE	DID+FE	DID+FE	DID+FE	DID+FE
Authors	382	212	382	212	382	212
Observations	764	424	764	424	764	424
R^2	0.324	0.261	0.018	0.038	0.021	0.015

Note: *** p<0.01, ** p<0.05, * p<0.1.

Table 6. Robustness Checks with Multi-ticker Articles

	Log(ArticleQuantity)		*Log(ArticleLength)*		*Log(TickerCoverage)*		*Log(SectorCoverage)*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment × *After*	1.006***	0.738***	0.025	0.063	0.849***	0.713***	0.242***	0.116**
	(0.106)	(0.135)	(0.043)	(0.058)	(0.119)	(0.149)	(0.047)	(0.059)
After	-1.084***	-0.827***	0.102***	0.013	-0.838***	-0.786***	-0.295***	-0.227***
	(0.094)	(0.133)	(0.031)	(0.053)	(0.094)	(0.147)	(0.036)	(0.055)
Log(ActivePeriod)	0.385***	0.417***	0.022	-0.016	0.469***	0.458***	0.058**	0.056
	(0.068)	(0.100)	(0.026)	(0.037)	(0.078)	(0.111)	(0.026)	(0.040)
Log(DaysOnSA)	0.482***	0.349**	0.006	0.082	0.338***	0.319*	0.243***	0.242***
	(0.136)	(0.172)	(0.042)	(0.057)	(0.130)	(0.174)	(0.044)	(0.060)
Specification	DID+FE	DID+FE	DID+FE	DID+FE	DID+FE	DID+FE	DID+FE	DID+FE
Authors	382	212	382	212	382	212	382	212
Observations	764	424	764	424	764	424	764	424
R^2	0.557	0.467	0.098	0.077	0.481	0.437	0.315	0.293

Note: *** p<0.01, ** p<0.05, * p<0.1.

Table 7. Summary Statistics: New Contributors

Variable	#Obs	Mean	Std. Dev.	Min	Median	Max
			Panel A: New Regular Contributors			
ArticleQuantity	318	2.53	5.26	1	1	47
ArticleQuality	318	0.52	0.45	0	0.50	1
CommentsPerArticle	318	15.96	24.01	0	8.50	226
ArticleQualityVariation	318	0.105	0.199	0	0	0.5
ArticleLength	318	1,079.53	767.66	176	882.50	6,222
TickerCoverage	318	2.12	4.46	1	1	45
SectorCoverage	318	1.35	1.00	1	1	8
ActivePeriod	318	546.92	196.33	20	628	711
			Panel B: New Premium Contributors			
ArticleQuantity	727	7.36	22.53	1	2	437
ArticleQuality	727	0.49	0.37	0	0.50	1
CommentsPerArticle	727	19.58	34.76	0	11.26	644.50
ArticleQualityVariation	727	0.240	0.239	0	0.37	0.5
ArticleLength	727	918.94	455.01	94	814	4,031
TickerCoverage	727	5.00	12.71	1	2	208
SectorCoverage	727	2.15	1.75	1	1	9
ActivePeriod	727	618.78	147.66	40	711	711

Table 8. Impact of Monetary Incentive for New Contributors

	Log(Article Quantity)	*Article Quality*	*Log(Comments PerArticle)*	*ArticleQuality Variation*	*Log(Article Length)*	*Log(Ticker Coverage)*	*Log(Sector Coverage)*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treatment	0.326***	-0.022	0.357***	0.093***	-0.060	0.276***	0.166***
	(0.043)	(0.029)	(0.078)	(0.014)	(0.038)	(0.038)	(0.022)
Log(ActivePeriod)	0.180***	-0.007	-0.051**	0.048***	-0.026***	0.141***	0.076***
	(0.011)	(0.008)	(0.021)	(0.003)	(0.010)	(0.010)	(0.005)
Authors	1,045	1,045	1,045	1,045	1,045	1,045	1,045
Observations	1,045	1,045	1,045	1,045	1,045	1,045	1,045
R^2	0.198	0.002	0.024	0.191	0.014	0.169	0.165

Note: *** p<0.01, ** p<0.05, * p<0.1.