



Many Destinations, Many Pathways: A Quantitative Analysis of Legitimate Peripheral Participation in Scratch

RUIJIA CHENG, University of Washington, USA

BENJAMIN MAKO HILL, University of Washington, USA

Although informal online learning communities have proliferated over the last two decades, a fundamental question remains: What are the users of these communities expected to learn? Guided by the work of Etienne Wenger on communities of practice, we identify three distinct types of learning goals common to online informal learning communities: the development of domain skills, the development of identity as a community member, and the development of community-specific values and practices. Given these goals, what is the best way to support learning? Drawing from previous research in social computing, we ask how different types of legitimate peripheral participation by newcomers—contribution to core tasks, engagement with practice proxies, social bonding, and feedback exchange—may be associated with these three learning goals. Using data from the Scratch online community, we conduct a quantitative analysis to explore these questions. Our study contributes both theoretical insights and empirical evidence on how different types of learning occur in informal online environments.

CCS Concepts: • **Human-centered computing** → **Empirical studies in collaborative and social computing**.

Additional Key Words and Phrases: online community, informal learning, newcomer participation

ACM Reference Format:

Ruijia Cheng and Benjamin Mako Hill. 2022. Many Destinations, Many Pathways: A Quantitative Analysis of Legitimate Peripheral Participation in Scratch. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 381 (November 2022), 26 pages. <https://doi.org/10.1145/3555106>

1 INTRODUCTION

In the last two decades, we have witnessed the proliferation of online communities that seek to promote informal learning through unstructured activities and interactions with others. For example, social computing scholars have documented the way users work together to learn creative and technical animation skills in communities such as *NewGrounds* [34], writing and web development skills in fan communities such as *FanFiction.net* and *Archive of Our Own (AO3)* [8, 15], and programming skills in creative coding communities such as *Scratch* [45]. These communities do not offer fixed lesson plans and offer little in the way of formal instruction. They are interest-driven, open to all, and free to join. Members learn at their own pace while pursuing their personal passions.

Skeptics of informal learning have pointed out that many users of such systems spend a substantial portion of their time on these websites commenting and socializing. How effective can a programming community be in supporting learning if its users constantly chat with each other while spending relatively little time writing or reading code? Proponents have responded that

Authors' addresses: Ruijia Cheng, rcheng6@uw.edu, University of Washington, Seattle, Washington, USA, 98195; Benjamin Mako Hill, University of Washington, Seattle, Washington, USA, 98195, makohill@uw.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2573-0142/2022/11-ART381

<https://doi.org/10.1145/3555106>

socialization promotes continued participation, a prerequisite to learning in informal settings where users are always at risk of not returning, and allows a collaborative process of learning and mentoring [8, 13, 49]. Others have suggested that learning to socialize is an important learning outcome itself [5, 14, 29]. An effective coding community could help its members develop important social skills in addition to technical ones [4].

There remains a deep disagreement within the social computing community about how learning communities should be designed to maximize learning. We believe that resolving these conversations is difficult because researchers often do not fully answer the following questions: What do we expect members of informal learning communities to learn? What types of behavior should designers encourage and support for each learning goal?

Our work begins to explore answers to these questions. Guided by Lave and Wenger's [33] influential work on communities of practice (CoPs) and Wenger et al.'s [53] more recent work on types of learning in CoPs, we describe three distinct categories of learning outcomes common to informal online learning communities: development of *domain* skills, development of *community* identity, and development of community *practices*. Using these three outcomes, we explore how different types of legitimate peripheral participation (LPP) by newcomers—contribution to core tasks, engagement with practice proxies, social bonding, and feedback exchange—are associated with these different learning outcomes using data from the Scratch online community [25]. We find that different types of participation are associated with learning outcomes differently. For example, making original projects as a newcomer is a good predictor of learning new computational concepts in the long term, while it is negatively associated with the formation of community identity and development of community practices.

Our paper makes several contributions. First, we believe that ours is the first study to provide a quantitative analysis of LPP and learning outcomes in a CoP. In doing so, our work suggests that users' early participation in an online community is associated with long-term learning outcomes—but the successful types of participation associated with different types of learning outcomes will differ. Second, we make an empirical contribution by analyzing 3 years of longitudinal data from Scratch. We believe that our work provides a template for future work that seeks to quantitatively investigate CoP theory in social computing in the context of a computational learning community. Finally, our work informs the design of new systems by pointing out potential ways to facilitate different newcomer participation patterns supporting a range of learning outcomes.

2 BACKGROUND

2.1 Communities of practice

Over the last three decades, Jean Lave and Etienne Wenger's ideas around CoP have been among the most important theories used by social computing scholars to understand how learning happens in communities. The term *community of practice* has been widely used in social computing scholarship to describe groups of people learning from each other while working toward a common interest or goal [e.g., 19, 31, 36, 50]. First introduced by Lave and Wenger [33] in 1991, CoP theory is grounded in ethnographic observation of apprenticeship relationships in communities of Liberian tailors, Mayan midwives, US Navy quartermasters, nondrinking alcoholics, and US supermarket meat cutters.

Learning in CoPs is described as occurring through LPP, the phenomenon through which newcomers begin to participate in a group by helping out with tasks that are easy and low-risk but still valuable and important. For example, a new midwife may begin by boiling water and cleaning scissors for other midwives. Through observing more experienced members performing

more complex and higher-stake tasks—and by practicing themselves in various ways—novices move from the periphery to more central roles.

2.2 Applications of CoP theory in social computing research

Although created to explain learning through traditional apprenticeships offline, the HCI and CSCW communities have embraced the CoP framework. Today, it is one of the most influential and highly cited theories related to learning in social computing. CoP has been widely adopted by social computing scholars because it is a good match for informal and individualized forms of learning that occur in many online settings.

Despite being cited hundreds of times in social computing scholarship—and tens of thousands of times in general—there have been very few attempts to explore or test CoP theory quantitatively. We are not aware of any such attempts in social computing scholarship. Furthermore, the work of applying CoPs to online settings, where groups are often extremely porous and diffuse, has required new theoretical work and argumentation. For example, Gruzd et al. [20] have argued that a single social media user can form a personal or “imaginary” CoP with other users in their network. Using a similar line of thinking, data scientists on Twitter who engage with the hashtag “#TidyTuesday” have been theorized to constitute a CoP because they share context, common interests, and a process of collaborative knowledge advancement and because they ask questions and interact with more experienced users [50].

A wide range of other examples of online CoPs that have been identified and studied in social computing includes design professionals on online critique platforms [31, 36], a Facebook group of Airbnb hosts who learn about new features and hosting strategies from each other [26], and fan fiction authors who gather online to develop and maintain the fan fiction website Archive of Our Own (AO3) [15]. In these examples and others, online CoPs are generally described as occurring in *affinity spaces* where members with similar interests and identities contribute to a shared collection of knowledge in a distributed manner [18]. An example of an online affinity space that is frequently described as a CoP is the Scratch online community—the empirical context of our study [39].

2.3 Types of learning in CoPs

One limitation in Lave and Wenger’s initial account is that it is vague about what exactly is being learned in a given CoP. This is an important omission because CoP members typically learn a range of different things as they become more experienced. For example, becoming a Mayan midwife in the Yucatan Peninsula involves much more than learning technical midwifery skills. It also involves learning about the Mayan midwifery community and the specific norms and values that shape midwifery in the Yucatan. In later work, Wenger et al. [53] attempted to address this omission by identifying three types of learning in CoPs: learning about *domain*, *community*, and *practice*. We present each type of learning in the following paragraphs.

First, learning about a *domain* refers to the acquisition of knowledge and skills necessary for a person to carry out the core tasks at the heart of a CoP. Many scholars of online communities are interested in how communities use LPP to support learning about some domain of knowledge and domain learning is often assumed to be the only learning goal in a CoP. For midwives, learning about a domain involves gaining knowledge related to successfully delivering infants. Depending on the community, domain knowledge may involve skills related to computer programming [45, 51], fan fiction writing [8], encyclopedia article editing [22], and so on. In the specific context of computational learning communities such as Scratch, domain learning means learning computational concepts and related programming skills [45].

The second type of learning involves the development of identity as a member of the *community*. This involves developing relationships, affinities, and a sense of belonging. As learners are accepted

by older members of the community, they gradually form an identity as a member of the community. The development of these relationships and affinities would play out similarly in various settings and typically involves the knowledge of other community members and the development of a sense of membership and commitment to the community [38].

Third and finally, learning a *practice* means assimilating “cultural artifacts, norms, and values” developed in the community over time [1]. By moving from peripheral to central forms of participation, learners develop by adjusting their social, work, and contribution style to match what is accepted and appreciated by community members. The specific reasons why one would be seen as accepted and appreciated in a CoP are very community-specific. However, indicators of practice development will typically involve expressions of appreciation and respect from others. For example, for midwives, it might involve the authority to lead or manage other midwives [33]. On Fanfiction.net, signs of practice development may involve high ratings and positive comments [8]. On Scratch, it might involve “loves” (i.e., likes) given by other users [3].

2.4 How does LPP promote different types of learning in CoPs?

Just as CoPs promote multiple types of learning, they also allow multiple types of participation. Although many studies of online activity reduce behavior to unidimensional concepts like “engagement,” there are many types of LPP that occur in CoPs. To identify different types of LPP relevant to online communities, we conducted a detailed search of the literature on CoPs in CSCW and social computing venues. We did not find any previous attempt to enumerate or classify different types of LPP in social computing. However, through our reading of the literature, we were able to identify four distinct types of LPP that are frequently discussed: contribution to core tasks, engagement with practice proxies (an important type of activity in CoP theory, which we will explain in detail below), the formation of social bonds, and feedback exchange. We do not claim that these four types are strictly mutually exclusive; nor do we claim that they form a comprehensive list. We merely offer them as examples of four different types of participation that capture some of the diversity of LPP. We discuss the limitations of our work based on our necessarily arbitrary identification of the types of LPP in §8.

The first type of LPP we identified, *contribution to core tasks*, refers to the work of newcomers toward a community’s explicit goal. Examples of contributions to core tasks include editing articles on Wikipedia, submitting code in an open source software project, and creating programming projects on Scratch. The second type of LPP is engagement with *practice proxies*. The term practice proxies refers to activities that allow newcomers to observe and participate in the socially salient aspects of others’ unfolding work practices [40]. For example, in CoPs for online citizen science, new users can learn to contribute by engaging with project documentation left by others [40]. The third type of LPP is *social bonding*. In Lave and Wenger’s [33] study of an alcoholics social group, the most important activities involved members forming interpersonal bonds and becoming friends. In online CoPs, social bonding is often enabled by social media features such as friending and following. The fourth form of LPP is participation in *feedback exchange*. Feedback exchange in online CoPs often takes place through public commenting on artifacts shared by community members.

Because participants in online groups learn different types of things, it stands to reason that the most effective forms of LPP for the promotion of learning might vary depending on the type of learning. In other words, not all forms of peripheral participation will be equally likely to help a newcomer develop in terms of every desirable learning outcome. For example, peripheral participation in socializing activities might help newcomers build knowledge about the community, but might not contribute to their knowledge of domain skills. Consequently, we might not expect that socializing in a coding-focused CoP would necessarily make one a better programmer. Similarly,

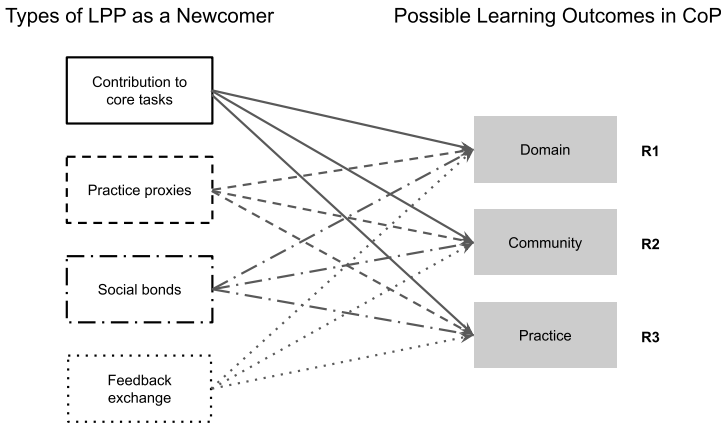


Fig. 1. Our research examines which type of newcomers’ LPP is associated with which learning outcome in CoP.

contribution to the core task of programming may not help to build social knowledge about the community and its members. Indeed, given limited time and resources, support for one learning outcome might come at the expense of others. For example, a study of the Scratch community showed that socialization among members in the comment section of Scratch projects can sometimes drive discussion away from programming related topics [49].

Inspired by the original CoP literature that focused primarily on the experience of newcomers and their progression in the community [33], we specifically focus on the participation of newcomers and their long-term learning outcomes. Figure 1 includes the four types of newcomer LPP we have described on the left and Wenger et al.’s [53] three types of learning in CoPs on the right. We draw 12 left-to-right arrows that represent all possible direct pathways between our four types of LPP and Wenger et al.’s [53] three learning outcomes. In the following sections, we explain our research questions (R1 to R3) and related work on how different types of LPP contribute to the three learning outcomes.

2.4.1 Domain. The social computing literature suggests that domain knowledge learning can be influenced by different types of LPP, although the specific effects remain unclear. For example, many studies suggest that by contributing to core tasks, newcomers can learn technical details and gain confidence. For example, new Wikipedia editors often start by making edits on topics that they are familiar with [6]. In the context of computational learning, novice programmers involved in open source projects often begin by reporting bugs and requesting features [51] and fanfiction authors learning to program contribute to the development of a fanfiction website by engineering very small features [15]. However, some also argue that contributing to core tasks can be stressful and overwhelming to newcomers, especially for those without much previous exposure to the domain.

Alternatively, newcomers can learn domain knowledge and skills with the help of practice proxies. For example, social media functionality that allows community members to view and download artifacts shared by others has been described as a method by which practice proxies can build domain knowledge among new designers [36]. Similarly, new Wikipedia editors rely on practice proxies by reading existing pages and page histories before making their initial contributions to the encyclopedia as a way of learning where, when, and how to edit [6]. In the context of computational learning, the evidence of the relationship between engagement with practice proxies and domain learning is mixed. While research shows that Scratch users who create remixes of others' projects (a common practice proxy in computational learning) will learn more computational concepts as they continue to participate in the community [12], they tend to display less innovation and originality in the programs they make [24].

Furthermore, while many computational learning platforms support socialization and feedback exchange, how these activities contribute to computational learning remains unclear. To explore how different types of LPP contribute to domain knowledge learning in CoPs, we ask the research question **R1**: How do different types of LPP affect domain knowledge learning?

2.4.2 Community. Qualitative evidence has shown that *engagement with practice proxies*, forming *social bonds*, and participating in *feedback exchange* can help newcomers develop an identity as a member of their *community*. Engagement with practice proxies can not only show newcomers how to master domain skills, but also provide newcomers with a sense of belonging. Lave and Wenger [33] describe how novice meat cutters in a US supermarket who lacked practice proxies had a low rate of interaction with more experienced butchers and did not feel like they were participating in a community to learn skills. In the context of Scratch, many users engage in remixing with the sole purpose of broadcasting each other's work and being part of a social group [24]. In the online CoP of R data scientists on Twitter, a large portion of newcomers join to build connections with others in the R community [50]. Furthermore, social computing research has shown that newcomers with connections to others tend to remain active in the community longer [32].

Previous research has indicated that newcomers' engagement in feedback exchange can also help users develop as community members. Feedback exchange in online CoPs has been described as a form of distributed mentoring in which mentoring relationships are not bounded by time, geographic space, and differences in expertise [8]. Similarly, an experiment on reader engagement on Wikipedia has shown that newcomers showed higher long-term participation rates when asked to provide feedback on articles [22]. To obtain quantitative evidence for how different types of LPP support members to develop community identity, we explore the research question **R2**: How do different types of LPP affect the development of community identity?

2.4.3 Practice. Few papers in the social computing literature have directly studied how *practice* in a community—i.e., being able to understand and reproduce what is valued by CoP members—can be supported. Some qualitative studies have implied that feedback exchange could help members of online CoPs learn community practices. For example, in forums dedicated to online dating, users offer each other comments and advice on self-presentation to understand how others use dating sites and what behaviors are considered appropriate [37]. Similarly, new Airbnb hosts critique each other's drafts of hosting descriptions to recognize effective hosting styles [26] and young UX enthusiasts comment on others' posts to help each other learn how to make an appropriate networking message in the community [31]. To further explore if and how different types of LPP can help learners learn community practices differently, we ask the research question **R3**: How do different types of LPP affect the learning of community-specific values and practices?



Fig. 2. Historical interfaces of Scratch and social features at the time when data used in our analysis were collected. Images obtained from Hill and Monroy-Hernández [25].

3 EMPIRICAL SETTING: SCRATCH COMMUNITY

We conducted our study using data from the Scratch community. Scratch was designed to support young people in learning to program using the Scratch programming language—a visual block-based language designed for children to learn basic programming [45, 46]. Scratch was designed based on constructionist learning principles [42, 44]. Programming primitives are represented by visual blocks that control the behavior of graphical objects on the screen called sprites. As shown in Figure 2a, users can drag and drop blocks together to build programs.

Although the Scratch programming language was designed for children to use on their own computers, the language has been integrated into a vibrant online community since 2007. Users in this community can view others’ project and share their programming projects with others [39].¹ As of July 2021, the Scratch online community has over 73 million registered users and over 79 million shared projects that span a diverse range of genres and themes. The large majority of Scratch users are between the ages of 8 and 16 years and the average age for new contributors is approximately 12 years.²

We chose Scratch for our study because it has been the center of several influential studies on informal learning in social computing and has been described as a CoP or site for situated learning in a range of previous studies [12, 13, 23]. Beyond its importance in social computing scholarship, Scratch is one of the largest online communities for children learning to code. Critically for our analysis, Scratch supports a variety of learning outcomes and forms of participation that allow us to test all three of our research questions and test the hypotheses associated with each arrow

¹<https://scratch.mit.edu/>

²All statistics about Scratch community activity and users are taken from the public information at: <https://scratch.mit.edu/statistics/>

in Figure 1. Scratch users have access to practice proxies in the form of affordances that support users in viewing, editing, and building on others' code by downloading and remixing each other's projects [39]. Previous studies have found evidence that remixing others' projects is associated with measures of learning about programming [12]. Scratch also supports users to build social connections by "friending" other users to form social networks and a shared sense of community [3, 4]. Scratch users can participate in feedback exchange by commenting on each others' projects, as shown in Figure 2b. For example, previous research has revealed collaborative debugging activities in project comments [49]. Finally, users can demonstrate social support by "loving" (upvoting) and "favoriting" (bookmarking) projects of others. We describe how these features map to our measure in the following section.

4 DATA AND MEASURES

Our data are drawn from publicly available *Scratch Research Dataset* (SRD), which includes comprehensive public data from the first 5 years of activity in Scratch between 2007 and 2012 [25]. We further restrict our analysis to projects created between 2 July 2009 and 10 April 2012 because both programming and social affordances in the site were consistent during this period.³ Our analysis relies on the tables *projects*, *users*, *pcomments*, *gcomments*, *project_blocks*, *favorites*, *lovers*, *viewers*, and *friends* from the SRD. Although this dataset is almost 10 years old, limiting ourselves to data from this period means that we can be more sure that the data are generated consistently and that others can reproduce our analysis using public data. To further support reproducibility, we have made the complete analytical code used in our analysis available.⁴ As in many other large online communities, Scratch users frequently register and participate only briefly before becoming inactive [11]. Because we are interested in how users' activities as newcomers predict long-term outcomes, we created a user-level dataset with 121,149 users who were active for at least one day after their initial registration during our period of observation.

For all of our research questions, we attempt to measure activity during two periods: users' time spent as newcomers and their time as established users. Following a series of previous studies on Scratch, we define the "newcomer period" of Scratch users as consisting of the first 14 days after each user creates their account on Scratch. Like participants in most online communities, Scratch users' transition from newcomers to established members is fuzzy and variable. From our own participation in the Scratch community, we sense that nearly all users would be considered newcomers on their first several days, but most users would no longer be considered newcomers after more than a month of activity. We choose 14 days because it falls between this range and because it is embedded in one of Scratch's key platform policies in that "Scratcher" status is only granted to users who have been on the platform for at least 14 days.⁵ Within Scratch, users with Scratcher status are typically not considered newcomers by other community members and have access to more advanced programming features, such as the Scratch Cloud Variables [13]. Other social computing studies have also defined the first 14-day period as the newcomer period and use user behaviors in this period to predict subsequent activities [7]. Because this choice of 14 days is necessarily arbitrary, we also repeated our analyses with 2, 7, and 30 days as the length of the newcomer period in a series of robustness checks. All these analyses lead to similar results and conclusions. We include Table 1 to give a sense of the distribution of activity in the first 14 days in our sample. A pattern—common in online community data from many sources—is that our

³https://en.scratch-wiki.info/wiki/Scratch_1.4

⁴<https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/0VXEMB>

⁵<https://en.scratch-wiki.info/wiki/Scratcher>

measures are heavily right-skewed. Although some users are very active, the median number for most measures is 0.

Measure	Mean	St. Dev.	Median	Range
Original projects made in the first 14 days	2.200	4.515	1	[0, 222]
Remixes made in the first 14 days	0.656	2.219	0	[0, 130]
Friends made in the first 14 days	1.269	9.115	0	[0, 1,487]
Comments made in the first 14 days	4.516	26.640	0	[0, 2,450]
Friends received in the first 14 days	0.832	2.877	0	[0, 206]
Loves received in the first 14 days	0.492	2.668	0	[0, 261]
Views received in the first 14 days	11.133	34.935	3	[0, 4, 161]
Remixes received in the first 14 days	0.111	1.528	0	[0, 446]
Comments received in the first 14 days	2.706	16.079	0	[0, 1,418]
Favorites received in the first 14 days	0.271	1.529	0	[0, 160]
CT concepts displayed in the first 14 days	2.886	2.364	3	[0, 6]
Total active duration (in days)	110.656	171.079	36	[1, 1,002]
Total new CT concepts	4.281	1.981	5	[0, 6]

Table 1. Descriptive statistics for a range of measures of user activity in Scratch including all activities that factor into our analysis.

Because most of the variation in the variables is between 0 and 1, we transform most of these variables into simple dichotomous measures for analysis. Each of the measures used in our regression analysis is shown in Table 2. Three variables—*Stayed?*, *New_CT_concepts?*, and *Received_new_loves?*—are our binary outcome variables. The rest are independent variables and controls. All of these are drawn directly from the SRD or computed through merging and aggregating user-level activities across tables in ways that are self-explanatory and well documented in our code.

The most important exception to this is our measures related to computational thinking (CT) concepts. CT is defined as “the thought processes involved in the formulation of problems and their solutions so that the solutions are represented in a form that can be carried out effectively by an information processing agent” [54]. Referring to Brennan and Resnick’s [2] interpretation of CT and categorization of CT concepts, Scratch’s designers designed programming blocks to embody specific CT concepts, including *loops*, *parallelism*, *events*, *conditionals*, *operations*, and *data* [45]. Our measure of *CT concepts* captures the number of each distinct concept that a user displays in their projects made in their first 14 days. The specific CT concept-to-block mapping that we use is drawn from the study by Dasgupta et al. [12] and can be found in our appendix in Table 4.

4.1 Measures

4.1.1 Dependent Variables. Because Scratch does not make any systematic attempt to measure learning, an analysis like ours must rely heavily on proxy measures of our outcomes. We constructed three outcome variables that correspond to our three research questions. The outcome variable for **R1** measures the *domain* being learned in Scratch—the computational skills. To construct a measure to capture the learning of computational skills, we follow the approach of previous quantitative studies on computational learning in Scratch. Our measure captures the size of the learners’ repertoire in terms of the number of types of CT concepts a user has demonstrated in their projects [12, 13, 48, 55]. We construct an outcome measure to capture the learning of computational

Variable	Mean	St. Dev.	Median	Range	Type
Independent variables:					
Made_original_project?	0.710	-	1	{0,1}	binary
Made_remix?	0.271	-	0	{0,1}	binary
Made_friend?	0.247	-	0	{0,1}	binary
Made_comment?	0.344	-	0	{0,1}	binary
Control variables:					
Has_been_friended?	0.271	-	0	{0,1}	binary
Has_been_loved?	0.170	-	0	{0,1}	binary
Has_been_viewed?	0.725	-	1	{0,1}	binary
Has_been_remixed?	0.057	-	0	{0,1}	binary
Has_been_commented?	0.329	-	0	{0,1}	binary
Has_been_favorited?	0.116	-	0	{0,1}	binary
CT_concepts	2.886	2.364	3	[0,6]	count
Outcome variables:					
Stayed?	0.678	-	1	{0,1}	binary
New_CT_concepts?	0.396	-	0	{0,1}	binary
Received_new_loves?	0.044	-	0	{0,1}	binary

Table 2. Distribution of variables used for regressions.

skills: *New_CT_concepts?*, a binary variable of whether the user uses any new CT concepts in their projects after the first 14 days.

For the outcome variable for **R2**, our proxy measure for *community* membership is the duration of a user staying active in the community. Although this is not a direct measure of learning, CoP theory suggests that users who are more integrated into a community will stay longer. We measure this outcome with a binary variable, *Stayed?*, that captures whether users will participate in any recorded community activities after 14 days. In this case, activity can include sharing original or remixed projects, posting comments, favoriting, and/or friending.

For **R3**, we want to test how participation in feedback exchange contributes to learning community-specific *practices* and values. Because it is difficult to directly measure users' learning of community values, we explore H3 using a proxy that seeks to measure the Scratch community's positive reaction to projects made by a user. Following previous work by Hill and Monroy-Hernández [24], we measure the reaction of the Scratch community to a user's projects as the number of loves received by that user. Specifically, we construct the measure for the outcome variable in R3 as *Received_new_love?*, a binary variable of whether the user received any loves on projects created after their first 14 days.

4.1.2 Independent Variables. We construct four independent variables to measure the four types of LPP that we describe in §2.4. First, we construct a measure for the LPP of contribution to core tasks. Since the Scratch community is designed to help children learn computer programming, the core

task in Scratch is to write code and share original projects [45]. Therefore, we construct a binary measure, *Made_original_project?*, that captures whether a user has shared an original project in their first 14 days. Second, we construct a measure for engagement with practice proxies. Drawing from previous literature on LPP in Scratch, we choose to operationalize practice proxies as *Made_remix?*, a binary variable of whether the user remixed in their first 14 days. Third, we construct a measure for social bonding. The most explicit way to form social bonds in Scratch is by adding other users as “friends.” In Scratch, friending means that a user has followed another user and will receive notifications when they post new projects.⁶ Because we want to measure whether a user is actively participating in forming social bonds, we construct a binary variable, *Made_friend?*, which captures whether a user has friended at least one other user in their first 14 days. Fourth, we constructed a measure for newcomers’ participation in feedback exchange. In Scratch, feedback is exchanged mainly in the form of comments [49]. There are two types of comments in Scratch: comments on projects and on “galleries.” We sum these two types of comments into the total number of comments posted and received by each user. Because we want to ensure that we are measuring whether a user is actively engaged in feedback exchange, we construct a binary variable, *Made_comment?*, which captures whether a user has posted a comment in their first 14 days.

4.2 Control variables

Based on previous work, we add a series of control variables that capture reasons that users might differ in their learning outcomes beyond differences in the four types of LPP we identify. For example, because previous literature has suggested that social support can affect newcomers’ learning [7, 17, 32], we control for incoming social approval on projects, constructed as four binary variables—*Has_been_loved?*, *Has_been_favorited?*, *Has_been_viewed?*, and *Has_been_remixed?*—which measure whether a user received loves, favorites, views, and remixes on their projects, respectively, in their first 14 days on Scratch. We also construct two binary variables, *Has_been_commented?* and *Has_been_friended?*, to control for the effect of passively received feedback and social bonds. Because the programming experience of users before joining the Scratch community may affect users’ learning of CT concepts and the reception of their projects within Scratch, we also included a control variable, *CT_concepts* that measures the total number of unique CT concepts used by a user in their first 14 days on Scratch. This control is only relevant in R1 and R3.

5 ANALYTIC PLAN

To answer R1 using the binary outcome variable *New_CT_concepts?*, we fit a logistic regression model on the dataset of 121,149 users using the GLM function in R.⁷ Because the distribution of the count control variable *CT_concepts* is right-skewed (i.e., most users used 0 CT concepts), we use a started log transformation in all the models involving this control (i.e., $\log(x + 1)$). Because our outcomes and all other measures are binary, our model is nonparametric, except for our *CT_concepts* variable. Our formal model for (M1) is as follows:

⁶<https://en.scratch-wiki.info/wiki/Friend>

⁷<https://www.rdocumentation.org/packages/stats/versions/3.6.2/topics/glm>

$$\log \left(\frac{\hat{p}(\text{New_CT_concepts?})}{1 - \hat{p}(\text{New_CT_concepts?})} \right) = \beta_0 + \beta_1 \text{Made_original_project?} + \beta_2 \text{Made_remix?} \\ + \beta_3 \text{Made_friend?} + \beta_4 \text{Made_comment?} + \beta_5 \text{Has_been_friended?} \\ + \beta_6 \text{Has_been_loved?} + \beta_7 \text{Has_been_viewed?} \\ + \beta_8 \text{Has_been_remixed?} + \beta_9 \text{Has_been_commented?} \\ + \beta_{10} \text{Has_been_favorited?} + \beta_{11} \log(\text{CT_concepts} + 1) \quad (1)$$

To answer R2 with the binary outcome variable *Stayed?*, we fit a logistic regression model (M2) that is identical to the model in Equation 1, except that it excludes the logarithmic-transformed variable *CT_concepts* and the associated parameter β_{11} . The dependent variable is as follows:

$$\log \left(\frac{\hat{p}(\text{Stayed?})}{1 - \hat{p}(\text{Stayed?})} \right) \quad (2)$$

To answer R3 with the binary outcome variable *Received_new_loves?*, we fit a logistic regression model on the full dataset. Our formal model (M3) is the same as that in Equation 1 but with the following dependent variable:

$$\log \left(\frac{\hat{p}(\text{Received_new_loves?})}{1 - \hat{p}(\text{Received_new_loves?})} \right) \quad (3)$$

6 RESULTS

As the first step in our analysis, we construct a series of bivariate graphs to show differences in the distributions of our outcomes across strata representing our key independent variables. These plots are shown in Figure 3. In general, we find large differences in answers to all three of our research questions. Of course, these plots are exploratory; the observed relationships could be driven by any number of confounders. As a result, we explore our research questions using the full array of controls in our regressions.

The regression results for all three research questions are shown in Table 3. Because interpreting marginal effects from regression models with many covariates can be challenging, especially in the case of GLM models such as logistic regression, we present a series of plots of model-predicted values in Figure 4. We do so by identifying the median values for each of our control variables (shown in Table 2) and then generating model predicted values for two prototypical users that vary only in terms of our key independent variables—e.g., a user who has (or has not) remixed, and so on. Each of the key independent variables not being visualized is also held at the sample medians.

6.1 Domain

For R1, we found that making original projects as a newcomer is positively associated with the usage of new computational concepts in the long term, whereas remixing as a newcomer has the opposite effect. As shown in Table 3, the odds that a user who created an original project in their first 14 days on Scratch would use new CT concepts afterward are 1.55 times the odds that an otherwise similar user who had not created an original project would do so ($\beta = 0.44$; SE = 0.03; $p < 0.001$). As seen in Figure 4 (R1a), our model predicts that approximately 81% of the prototypical users who had not made original projects in their first 14 days would use new CT concepts afterward, whereas approximately 87% of otherwise similar users who had made original projects would. Furthermore, we find that the odds that a user who remixed in their first 14 days would use a new CT concept afterward are 0.74 times the odds of an otherwise similar user who had not done so ($\beta = -0.30$; SE = 0.02; $p < 0.001$). As seen in Figure 4 (R1b), our model predicts that approximately 81% of prototypical users who had not remixed in their first 14 days would use

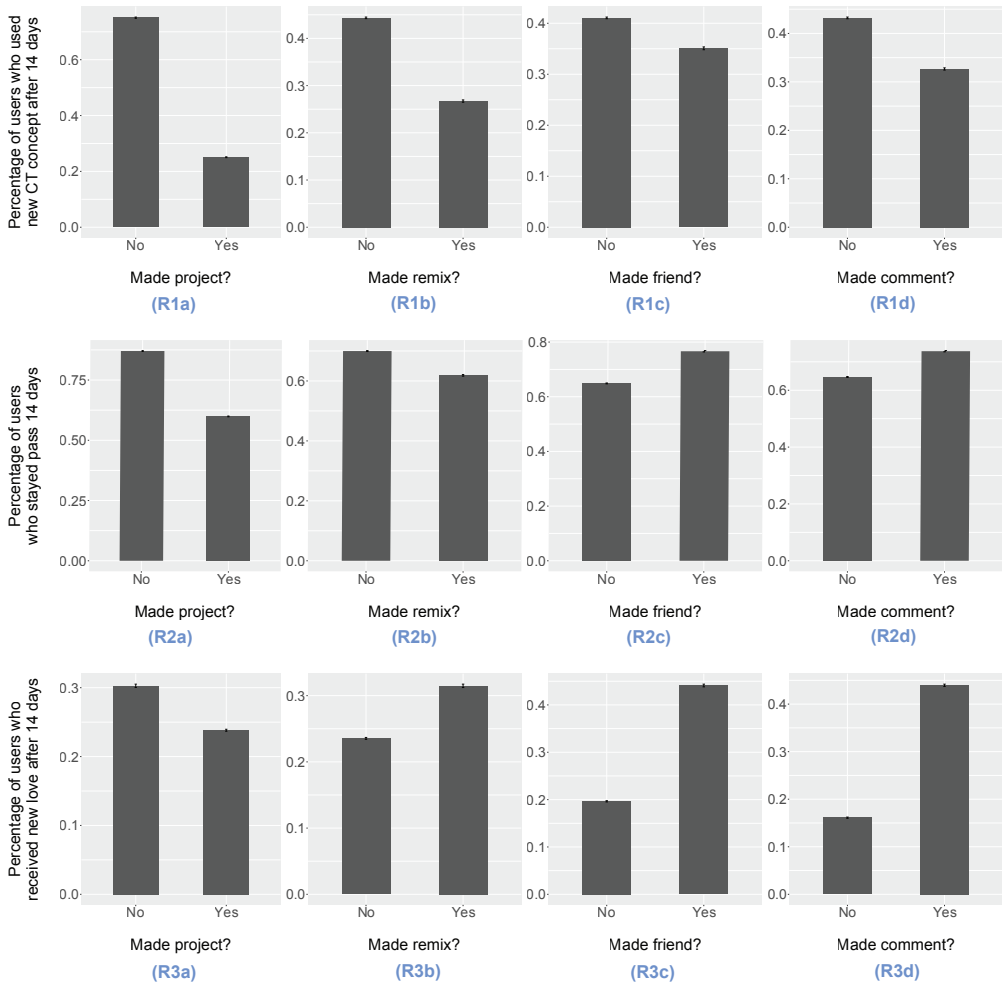


Fig. 3. Bivariate plots that show the differences in the distributions of outcome variables across strata of our dataset that reflect the independent variables associated with each of our research questions.

new CT concepts afterward, whereas approximately 76% of the otherwise similar users who had remixed would. The effects of friending and commenting are not statistically significant.

6.2 Community

For R2, we found that making original projects and remixing as a newcomer are negatively associated with long-term stay in the community, whereas making friends and commenting as a newcomer were positively correlated with the probability of staying in the community for 14 days. The odds that a user who created an original project in their first 14 days in Scratch would stay in the community afterward are 0.53 times the odds that an otherwise similar user who had not created an original project would do so ($\beta = -0.63$; $SE = 0.02$; $p < 0.001$). As seen in Figure 4 (R2a), our model predicts that approximately 91% of prototypical users who had not made original projects in the first 14 days will stay longer than 14 days, whereas only approximately 84% of otherwise

	<i>New CT concept?</i> (M1)	<i>Stayed?</i> (M2)	<i>Received new love?</i> (M3)
(Intercept)	1.47*** (0.01)	2.27*** (0.02)	-1.16*** (0.01)
<i>Made_original_project?</i>	0.44*** (0.03)	-0.63*** (0.02)	-0.71*** (0.04)
<i>Made_remix?</i>	-0.30*** (0.02)	-0.24*** (0.02)	0.00 (0.02)
<i>Made_friend?</i>	0.01 (0.02)	0.34*** (0.02)	0.28*** (0.02)
<i>Made_comment?</i>	0.00 (0.02)	0.48*** (0.02)	0.95*** (0.02)
<i>Has_been_friended?</i>	0.20*** (0.02)	0.38*** (0.02)	0.42*** (0.02)
<i>Has_been_loved?</i>	-0.02 (0.03)	0.19*** (0.02)	0.70*** (0.02)
<i>Has_been_viewed?</i>	-1.25*** (0.02)	-1.70*** (0.03)	-0.76*** (0.03)
<i>Has_been_remixed?</i>	0.04 (0.03)	0.12*** (0.03)	0.31*** (0.03)
<i>Has_been_commented?</i>	-0.13*** (0.02)	0.00 (0.02)	0.26*** (0.02)
<i>Has_been_favorited?</i>	-0.04 (0.03)	0.17*** (0.03)	0.41*** (0.02)
log1p(<i>CT_concepts</i>)	-1.19*** (0.02)	N/A	0.17*** (0.02)
AIC	125313.82	133601.48	119286.65
BIC	125430.28	133708.23	119403.11
Log Likelihood	-62644.91	-66789.74	-59631.33
Deviance	125289.82	133579.48	119262.65
McFadden's pseudo R-squared	0.23	0.12	0.14
Num. obs.	121149	121149	121149

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Table 3. Logistic regression models that predict the probability that a user uses a new CT concept (M1), stays active (M2), and receives loves on new projects (M3) created after their first 14 days on the community. The models are fit to the user-level dataset that includes aggregated activities of 121, 149 users.

similar users who had made original projects would stay. In terms of the effect of remixing, the odds that users who remixed others' projects in their first 14 days will stay longer than 14 days are 0.79 times the odds of a user who did not do so ($\beta = -0.24$; $SE = 0.01$; $p < 0.001$). As seen in Figure 4 (R2b), our model predicts that more than 90% of prototypical users who did not remix in their first 14 days would stay in the community beyond 14 days, whereas approximately 88% of otherwise similar users who had remixed would. In contrast, the odds that users who friended others in their first 14 days would stay past that time period are 1.40 times the odds that a user who did not friend anyone would ($\beta = 0.34$; $SE = 0.02$; $p < 0.001$). Figure 4 (R2c) shows that our model predicts that approximately 90% of our prototypical users who had not friended anyone in their first 14 days




Fig. 4. Plots of model predicted probabilities corresponding to each of our research questions. Each figure shows the model predicted probabilities for two prototypical users who have median values for each of our control variables; these users vary only in terms of the key independent variable in the corresponding hypothesis. Error bars reflect the marginal effects of our key independent variable ($1.96 \times SE$).

would stay in the community past 14 days, whereas approximately 93% of otherwise similar users who had friended people would do so. For users who posted feedback in the community in their first 14 days, their odds of staying in Scratch after 14 days are 1.62 times the odds of users who did not ($\beta = 0.48$; $SE = 0.01$; $p < 0.001$). Figure 4 (R2d) shows that our model predicts that only approximately 90% of prototypical users who had not posted any comments in their first 14 days would stay in the community past 14 days, whereas approximately 94% of otherwise similar users who had posted comments would.

6.3 Practice

For R3, we found that while making original projects as a newcomer is negatively associated with receiving loves from the community in the long term, making friends and commenting as a newcomer have a positive relationship with the outcome. Our model predicts that users who created an original project in their first 14 days in Scratch have 0.49 times the odds of receiving loves for their projects after 14 days compared with those who did not post feedback in the same newcomer period ($\beta = -0.71$; $SE = 0.04$; $p < 0.001$). As seen in Figure 4 (R3a), our model predicts that approximately 24% of prototypical users who had not made original projects in their first 14 days will stay longer than 14 days, whereas only approximately 13% of otherwise similar users who had made original projects would. In contrast, the odds that users who friended others in their first 14 days would receive new loves are 1.32 times the odds that a user who did not friend anyone would do so ($\beta = 0.28$; $SE = 0.02$; $p < 0.001$). Figure 4 (R3c) shows that our model predicts that approximately 24% of our prototypical users who had not friended anyone in the first 14 days would stay in the community past 14 days, whereas approximately 29% of otherwise similar users who had friended people would do so. In terms of feedback exchange, we find evidence that users who posted feedback in the community in their first 14 days would have 2.58 times the odds to receive loves for their projects after 14 days compared to those who did not post feedback in their newcomer period ($\beta = 0.95$; $SE = 0.02$; $p < 0.001$). As shown in Figure 4 (R3d), the predicted probability of receiving loves for new projects created after 14 days is only approximately 25% for prototypical users who did not comment during their initial period, but this value is close to 45% for those who did. The effect of remixing is not statistically significant.

7 DISCUSSION

In this paper, we present a quantitative study with data from the Scratch community that investigates how newcomers' LPP contributes to different kinds of learning in CoP. Our paper makes a contribution to social computing theory by drawing in Wenger et al.'s [53] three types of learning from the broader CoP literature: development of domain skills, development of identity as a community member, and development of community-specific values and practices. In our synthesis of the CoP literature in social computing, we identified four types of LPP common among newcomers (contribution to core tasks, engagement with practice proxies, social bonding, and feedback exchange) and formed three research questions related to understanding how the early participation of users in the community will predict learning outcomes in CoPs.

We contribute what we think is the first quantitative test of CoP theory in the context of social computing. Analytically speaking, our work tested 12 hypotheses that correspond to the 12 possible relationships between our independent and dependent variables (i.e., the lines in Figure 1). Our work finds that most of these relationships are statistically significant, but that the signs and relative magnitudes of parameters associated with each type of LPP vary across the types of learning outcomes. At a very high level, our results provide concrete evidence that different types of LPP have different relationships with different important facets of learning in online CoPs. What is productive for some types of learning outcomes is unhelpful for others, and vice versa. We devote the remainder of our discussion to unpacking our specific results.

7.1 Supporting contribution to core tasks

Our study finds quantitative evidence for the belief that learners in a CoP will learn the domain by being part of the core tasks in the community. This finding is consistent with the constructionist design approach behind Scratch. It appears that by participating in epistemologically relevant

project creation, learners can appropriate the abstract concepts behind their hands-on experiences [42, 44, 45].

That being said, contribution to core tasks is negatively associated with our measure of learning as membership identity and community-specific practices. We speculate that because creating a programming project from scratch can be challenging for novices, some newcomers might feel discouraged and, as a result, turn away from the community. This is consistent with the results reported in previous literature, which indicate that it can be challenging for newcomers to follow community norms while making contributions [21]. Existing resources reflecting good practices are often composed by experienced users and contain contextual information that is not beginner-friendly [15, 36]. These might be obstacles for newcomers looking to adopt community values.

Designers of the Scratch language have made enormous efforts to lower the bar of programming for novices [45]. Still, making a new project, however simple, might be overwhelming for someone without programming experience. Future design of informal learning communities like Scratch could consider offering easier options for newcomers to engage in core tasks. For example, in some other online CoPs dedicated to software development, newcomers can participate in simple, collaborative, programming-related tasks such as reporting bugs [51]. Scratch might try to direct newcomers toward types of newcomer-friendly programming tasks such as building on example templates, identifying bugs in collaborative projects, and helping more experienced users with a small part of their projects.

7.2 Supporting engagement with practice proxies

Our exploration shows that engagement with practice proxies is negatively associated with the development of domain knowledge and the formation of membership identity. Although previous work on Scratch shows that users who remixed projects containing a CT concept tended to use the same CT concept in their own projects [12], our findings suggest that this may not contribute to the expansion of computational knowledge and skills in general. Our results are consistent with previous findings on remixing, in which users tend to show less innovation and originality when building from the most frequently remixed programs [24].

More research needs to be conducted to understand exactly why remixing negatively impacts some dimensions of learning. One possible explanation is that newcomers may have trouble understanding specific choices in projects made by more experienced members and may not have the confidence needed to try and build upon them. In Kaggle, novices face a similar challenge of understanding and reusing publicly shared code by experts due to incomplete documentation on dependencies, missing rationales, and lack of context [10]. In these cases, simply providing newcomers with practice proxies is likely not enough.

Possible solutions might include additional scaffolding by established community members to document their steps and reasoning processes in the first place. Wang et al. [52] and Rule et al. [47] provide great guidelines and examples for such design in the context of computational learning. Informal learning communities should also offer newcomers the opportunity to ask questions to experienced members. Such a design should scaffold the asker to clearly communicate which step they are confused about in context and support answerers to effectively showcase their procedures and rationales. For example, in the context of Scratch, a synchronous Q&A function (e.g., an opt-in chat room similar to the one proposed by Ford et al. [16]) could be implemented in the remix feature so that newcomers could discuss the coding mechanisms with the original creators of the projects being remixed.

7.3 Supporting newcomer's socialization

Newcomers' social bonding activity is associated with the formation of community identity. Such socialization activities are also positively correlated with the development of community-specific practices. This finding can be read as a strong response to the skeptics of informal learning in online CoPs who argue that chatting and socialization will distract users. Our results complement other work that shows that socialization can be a legitimate pathway to learning what the community values [9]. One possible interpretation is that social bonding can facilitate computational participation [30]—learning through participating in socially situated contexts.

Better ways to socialize newcomers is a perennial topic in social computing. Our results suggest that success in these efforts can support a variety of learning outcomes. Proactive efforts to reach out to new users with comments are a useful way to jumpstart social bonding processes. Designers of informal online learning systems may consider intentionally helping newcomers socialize with each other and older members through community events or algorithmic matching mechanisms, for example, making the transition from identity-based connections to bond-based connections possible [32].

7.4 Supporting feedback exchange

Feedback exchange is positively associated with our measures of learning about both community and practice. Complementing previous studies of collaborative debugging [49], our study offers quantitative evidence that socially oriented newcomer participation can contribute to community practices. Despite the promise of feedback exchange on learning, newcomers refrain from publicly offering feedback because they are self-conscious about their social status and are not confident in their expertise [36]. Nonpublic or anonymous feedback systems might be established so that newcomers do not have to reveal themselves while interacting with community members [16]. Alternatively, scaffolds for composing discussion messages could be introduced so that newcomers can have some idea about what to say in comments and participate in feedback exchange more easily [28, 35].

8 THREATS TO VALIDITY

A series of articles in the social computing literature have argued about whether deeply engaged and highly active users in online communities are “born and not made” [27, 41, 43]. In its analytic structure, our work has quite a bit in common with *born versus made* studies by Panciera et al. [41] and Huang et al. [27]. Both studies seek to make predictions about users' long-term behavior using measures of early activity. There are also important differences between these approaches and ours. Born/made studies typically look at users' very first engagement. Although we take a longer view in the construction of our measures, our results are not sensitive to this decision. Although we include a large body of control variables to address *ex ante* differences between users, we cannot fully rule out the explanation that users who are more active early on are systematically different from users who are not. In contrast, we might argue that our work points to an alternative explanation for early studies that suggest that users are “born and not made.” Both our theoretical framework and empirical results suggest that early activity may not be best thought of as an indicator of existing differences in knowledge, commitment, and skills, but rather as a pathway through which users can develop.

Furthermore, we define the “newcomer” period as the first 14 days after a user creates an account. Although this decision is informed by our experiences with Scratch and although robustness checks on user activities in the first 2, 7, and 30 days lead to consistent results, the choice of 14 days is arbitrary and fails to capture important nuance in the way that Scratch users understand their own

status, and the status of others, within their community. We hope that future work can improve on this very limited measurement approach.

Another limitation of our work is that the four types of LPP that we identified are not a comprehensive list of all possible types of LPP in online communities, either in general or in Scratch in particular. As the result of our review of relevant social computing literature, our enumeration of these four types is put forward with humility and knowledge that is limited and incomplete. Although we know that our list is not comprehensive, we hope that our enumeration and measurement of four distinct types of LPP are sufficient to convey our high-level argument that there are different types of newcomer engagement that can lead to different learning pathways. Future research could explore other types of LPP and their impact on long-term learning outcomes.

Our work is also limited in that our measures are proxies rather than direct measures of the concepts of interest. This is particularly important to acknowledge in terms of our learning outcome measures. We rely on proxies because we simply do not have a way to evaluate actual learning in Scratch. For example, we use differences in incoming loves as a proxy for an individual's learning about community-specific practices and values. It is important to remember that receiving positive community reactions is an outcome of learning community practices, not a direct measure of learning. Another example is that we use a measure of whether a user stays in the community for longer than 14 days to proxy for a learner's development of membership identity. We rely on proxies because sociopsychological concepts are difficult to directly measure in informal settings like Scratch, where there is little or no ability to conduct tests or even communicate with users. Although we draw heavily from other work that has developed and evaluated proxy measures in Scratch, all of these attempts are risky and prone to noise, or worse [12, 13]. Although far from perfect, we believe that these proxies provide valuable insight into our research questions.

While our research questions describe causal relations, our regression analyses can only provide correlational evidence. As is often the case in cross-sectional analyses, these relationships might be due to variables that are correlated with, but not caused by, our predictors. We hope that design-based field experiments can be conducted in the future to identify causal relationships between specific types of newcomer LPP and the specific kind of learning in CoPs. Until then, we offer our results as tentative evidence. Furthermore, our analysis focuses only on the main effect of the independent variables. In an earlier version of the analysis, we include two-way interactions of our independent variables as predictors in our models, and we include the results in Table 6 in the appendix. Since understanding the interaction relationships between our independent variables is not the goal of this analysis and it is difficult to interpret the coefficients of interactions, we did not integrate these models into our main analysis. We urge future researchers to build on our work and explore the interactions and combination of different types of LPP and their contribution to different types of learning.

Our work is also limited in that all our outcomes are dichotomous. This is a choice we made for two reasons. First, it means that our analysis is at a lower risk of violating parametric modeling assumptions. Second, doing so also captures the vast majority of the variation in our outcomes. However, this choice means that our results only paint part of the picture. Although we show that sharing new projects and remixing early on are associated with a higher chance of using CT concepts later, we do not know if more remixing leads to more new CT concepts among users who continue. Although we leave this question for future work, we include an exploratory analysis in Table 5 in our appendix that reproduces our three models with continuous measures of our dependent variables among the relatively small proportion of users for whom our outcomes were equal to 1. This is conceptually similar to a hurdle model. Although limited in several ways, these results suggest that our dichotomous measures, which are strong predictors of new CT concepts, staying active, and receiving loves, are not necessarily strong predictors of higher amounts of

learning among users who have learned at all. We hope to refine and continue to explore the relationship between the types of newcomer LPP and the magnitude of learning outcomes in future studies.

Finally, our analysis focuses on Scratch as the single empirical setting. While Scratch is large, active, and supports a wide range of activities, it is also unique in many aspects: its user population mostly consists of young people, the community focuses on computational learning, and it contains affordances such as remixing and viewing others' full code. We do not know if our findings from the Scratch online community will generalize to other settings or population groups. In addition, we analyzed activities during a single time period, which was early in the Scratch community's lifetime. Since the affordances of Scratch have changed and the community itself has also grown since then, we cannot know how well our findings will translate to Scratch today. We hope that our case study offers a framework for studying different types of LPP and learning in online CoPs. Future researchers could answer our research questions in other settings and compare their results with our findings.

9 CONCLUSION

In this study, we present a quantitative approach to engaging with CoP theory in an informal online learning context. Although previous studies in social computing have treated learning outcomes as a single dimension, we build on work by Wenger et al. [53] to describe three distinct types of learning that can be supported in online CoPs: learning about the domain, learning about the community, and learning about the practice. We also identify four forms of LPP that are common in informal online learning contexts. We use historical data from Scratch to answer a series of research questions about whether certain types of newcomer LPP contribute to certain kinds of learning. Our results suggest that there is a range of possible pathways to a range of distinct learning outcomes.

Taking a broad exploratory approach instead of a more narrow hypothesis testing structure, we find evidence of relationships between types of LPP and learning outcomes that we believe are entirely untheorized in the CoP literature in social computing. For example, we find that newcomer socializing is associated with learning about the community and its practices. Our study offers theoretical and empirical contributions to social computing research on informal learning settings, as well as practical implications on how to best design informal online learning systems. In summary, our study suggests that online communities afford many possible destinations in learning and many pathways to each.

ACKNOWLEDGMENTS

We would like to acknowledge members of the Community Data Science Collective for their feedback and support. We also want to thank our anonymous reviewers for their feedback. We want to thank the Scratch team—especially Mitchel Resnick, Natalie Rusk, and Andrés Monroy-Hernandez—for building the Scratch online communities and for making data about Scratch available that allowed us to conduct this work. Finally, we want to thank members of the Scratch online community who generated data for this research and who inspired our work.

REFERENCES

- [1] Sasha A Barab and Thomas Duffy. 2000. From practice fields to communities of practice. *Theoretical foundations of learning environments* 1, 1 (2000), 25–55.
- [2] Karen Brennan and Mitchel Resnick. 2012. New frameworks for studying and assessing the development of computational thinking. In *Proceedings of the 2012 annual meeting of the American educational research association, Vancouver, Canada*, Vol. 1. 25.
- [3] Karen Brennan and Mitchel Resnick. 2013. Imagining, creating, playing, sharing, reflecting: How online community supports young people as designers of interactive media. In *Emerging technologies for the classroom*. Springer, 253–268.
- [4] Karen Brennan, Amanda Valverde, Joe Prempeh, Ricarose Roque, and Michelle Chung. 2011. More than code: The significance of social interactions in young people’s development as interactive media creators. In *EdMedia+ Innovate Learning*. Association for the Advancement of Computing in Education (AACE), 2147–2156.
- [5] Amy Bruckman. 2005. *Learning in Online Communities*. Cambridge University Press, 461–472. <https://doi.org/10.1017/CBO9780511816833.028>
- [6] Susan L Bryant, Andrea Forte, and Amy Bruckman. 2005. Becoming Wikipedian: transformation of participation in a collaborative online encyclopedia. In *Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work*. 1–10.
- [7] Moira Burke, Cameron Marlow, and Thomas Lento. 2009. Feed me: motivating newcomer contribution in social network sites. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 945–954.
- [8] Julie Campbell, Cecilia Aragon, Katie Davis, Sarah Evans, Abigail Evans, and David Randall. 2016. Thousands of positive reviews: Distributed mentoring in online fan communities. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. 691–704.
- [9] Ruijia Cheng, Sayamindu Dasgupta, and Benjamin Mako Hill. 2022. How Interest-Driven Content Creation Constrains Opportunities for Informal Learning: A Case Study on Novices’ Use of Data Structures in Scratch. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems - CHI ’22*. ACM Press, New Orleans, LA, USA. <https://doi.org/10.1145/3491102.3502124>
- [10] Ruijia Cheng and Mark Zachry. 2020. Building Community Knowledge In Online Competitions: Motivation, Practices and Challenges. In *Proc. ACM Hum.-Comput. Interact.*, CSCW2, Vol. 4. <https://doi.org/10.1145/3415250>
- [11] Kevin Crowston and Isabelle Fagnot. 2018. Stages of motivation for contributing user-generated content: A theory and empirical test. *International Journal of Human-Computer Studies* 109 (2018), 89–101. <https://doi.org/10.1016/j.ijhcs.2017.08.005>
- [12] Sayamindu Dasgupta, William Hale, Andrés Monroy-Hernández, and Benjamin Mako Hill. 2016. Remixing as a pathway to computational thinking. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. 1438–1449.
- [13] Sayamindu Dasgupta and Benjamin Mako Hill. 2018. How “wide walls” can increase engagement: evidence from a natural experiment in Scratch. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [14] Sarah Evans, Katie Davis, Abigail Evans, Julie Ann Campbell, David P Randall, Kodlee Yin, and Cecilia Aragon. 2017. More than peer production: Fanfiction communities as sites of distributed mentoring. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. 259–272.
- [15] Casey Fiesler, Shannon Morrison, R. Benjamin Shapiro, and Amy S. Bruckman. 2017. Growing Their Own: Legitimate Peripheral Participation for Computational Learning in an Online Fandom Community. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (Portland, Oregon, USA) (CSCW ’17)*. Association for Computing Machinery, New York, NY, USA, 1375–1386. <https://doi.org/10.1145/2998181.2998210>
- [16] Denae Ford, Kristina Lustig, Jeremy Banks, and Chris Parnin. 2018. “We Don’t Do That Here”: How Collaborative Editing with Mentors Improves Engagement in Social Q&A Communities. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI ’18*. ACM Press, Montreal QC, Canada, 1–12. <https://doi.org/10.1145/3173574.3174182>

- [17] Dena Ford, Justin Smith, Philip J Guo, and Chris Parnin. 2016. Paradise unplugged: Identifying barriers for female participation on stack overflow. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. 846–857.
- [18] James Paul Gee. 2005. Semiotic social spaces and affinity spaces. *Beyond communities of practice language power and social context* 214232 (2005).
- [19] Sarah Gilbert. 2016. Learning in a Twitter-based community of practice: an exploration of knowledge exchange as a motivation for participation in# hcsmsca. *Information, Communication & Society* 19, 9 (2016), 1214–1232.
- [20] Anatoliy Gruzd, Barry Wellman, and Yuri Takhteyev. 2011. Imagining Twitter as an imagined community. *American Behavioral Scientist* 55, 10 (2011), 1294–1318.
- [21] Aaron Halfaker, R Stuart Geiger, Jonathan T Morgan, and John Riedl. 2013. The rise and decline of an open collaboration system: How Wikipedia’s reaction to popularity is causing its decline. *American Behavioral Scientist* 57, 5 (2013), 664–688.
- [22] Aaron Halfaker, Os Keyes, and Dario Taraborelli. 2013. Making peripheral participation legitimate: reader engagement experiments in wikipedia. In *Proceedings of the 2013 conference on Computer supported cooperative work*. 849–860.
- [23] Samantha Hautea, Sayamindu Dasgupta, and Benjamin Mako Hill. 2017. Youth perspectives on critical data literacies. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 919–930.
- [24] Benjamin Mako Hill and Andrés Monroy-Hernández. 2013. The cost of collaboration for code and art: Evidence from a remixing community. In *Proceedings of the 2013 conference on Computer supported cooperative work*. 1035–1046.
- [25] Benjamin Mako Hill and Andrés Monroy-Hernández. 2017. A Longitudinal Dataset of Five Years of Public Activity in the Scratch Online Community. *Scientific Data* 4 (Jan. 2017), 170002. <https://doi.org/10.1038/sdata.2017.2>
- [26] Maya Holikatti, Shagun Jhaver, and Neha Kumar. 2019. Learning to Airbnb by engaging in online communities of practice. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–19.
- [27] Shih-Wen Huang, Minhyang Suh, Benjamin Mako Hill, and Gary Hsieh. 2015. How activists are both born and made: An analysis of users on change. org. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 211–220.
- [28] Julie S Hui, Darren Gergle, and Elizabeth M Gerber. 2018. Introassist: A tool to support writing introductory help requests. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [29] Mizuko Ito. 2013. *Hanging out, messing around, and geeking out: Kids living and learning with new media*. The MIT Press.
- [30] Yasmin B. Kafai. 2016. From Computational Thinking to Computational Participation in K–12 Education. *Commun. ACM* 59, 8 (July 2016), 26–27. <https://doi.org/10.1145/2955114>
- [31] Yubo Kou, Colin M Gray, Austin L Toombs, and Robin S Adams. 2018. Understanding social roles in an online community of volatile practice: A study of user experience practitioners on reddit. *ACM Transactions on Social Computing* 1, 4 (2018), 1–22.
- [32] Robert E. Kraut, Paul Resnick, and Sara Kiesler. 2011. *Building Successful Online Communities: Evidence-Based Social Design*. MIT Press, Cambridge, Mass.
- [33] Jean Lave and Etienne Wenger. 1991. *Situated learning: Legitimate peripheral participation*. Cambridge university press.
- [34] Kurt Luther, Kevin Ziegler, Kelly E Caine, and Amy Bruckman. 2009. Predicting successful completion of online collaborative animation projects. In *Proceedings of the seventh ACM conference on Creativity and cognition*. 391–392.
- [35] Stephen MacNeil, Zijian Ding, Kexin Quan, Thomas j Parashos, Yajie Sun, and Steven P Dow. 2021. Framing Creative Work: Helping Novices Frame Better Problems through Interactive Scaffolding. In *Creativity and Cognition*. 1–10.
- [36] Jennifer Marlow and Laura Dabbish. 2014. From Rookie to All-Star: Professional Development in a Graphic Design Social Networking Site. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing - CSCW ’14*. ACM Press, Baltimore, Maryland, USA, 922–933. <https://doi.org/10.1145/2531602.2531651>
- [37] Christina Masden and W. Keith Edwards. 2015. Understanding the Role of Community in Online Dating. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (Seoul, Republic of Korea) (CHI ’15)*. Association for Computing Machinery, New York, NY, USA, 535–544. <https://doi.org/10.1145/2702123.2702417>
- [38] David W McMillan and David M Chavis. 1986. Sense of community: A definition and theory. *Journal of community psychology* 14, 1 (1986), 6–23.
- [39] Andrés Monroy-Hernández. 2007. ScratchR: Sharing User-Generated Programmable Media. In *Proceedings of the 6th International Conference on Interaction Design and Children (IDC ’07)*. ACM, New York, NY, 167–168. <https://doi.org/10.1145/1297277.1297315>
- [40] Gabriel Mugar, Carsten Østerlund, Katie DeVries Hassman, Kevin Crowston, and Corey Brian Jackson. 2014. Planet Hunters and Seafloor Explorers: Legitimate Peripheral Participation through Practice Proxies in Online Citizen Science. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (Baltimore, Maryland, USA) (CSCW ’14)*. Association for Computing Machinery, New York, NY, USA, 109–119. <https://doi.org/10.1145/2531602.2531651>

[//doi.org/10.1145/2531602.2531721](https://doi.org/10.1145/2531602.2531721)

- [41] Katherine Panciera, Aaron Halfaker, and Loren Terveen. 2009. Wikipedians are born, not made: a study of power editors on Wikipedia. In *Proceedings of the ACM 2009 international conference on Supporting group work*. 51–60.
- [42] Seymour A Papert. 1980. *Mindstorms: Children, computers, and powerful ideas*. Basic books.
- [43] Jennifer Preece and Ben Shneiderman. 2009. The reader-to-leader framework: Motivating technology-mediated social participation. *AIS transactions on human-computer interaction* 1, 1 (2009), 13–32.
- [44] Mitchel Resnick, Amy Bruckman, and Fred Martin. 1996. Pianos not stereos: Creating computational construction kits. *interactions* 3, 5 (1996), 40–50.
- [45] Mitchel Resnick, John Maloney, Andrés Monroy-Hernández, Natalie Rusk, Evelyn Eastmond, Karen Brennan, Amon Millner, Eric Rosenbaum, Jay Silver, Brian Silverman, et al. 2009. Scratch: programming for all. *Commun. ACM* 52, 11 (2009), 60–67.
- [46] Ricarose Roque, Yasmin Kafai, and Deborah Fields. 2012. From Tools to Communities: Designs to Support Online Creative Collaboration in Scratch. In *Proceedings of the 11th International Conference on Interaction Design and Children (Bremen, Germany) (IDC '12)*. Association for Computing Machinery, New York, NY, USA, 220–223. <https://doi.org/10.1145/2307096.2307130>
- [47] Adam Rule, Ian Drosos, Aurélien Tabard, and James D. Hollan. 2018. Aiding Collaborative Reuse of Computational Notebooks with Annotated Cell Folding. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 150 (Nov. 2018), 12 pages. <https://doi.org/10.1145/3274419>
- [48] Christopher Scaffidi and Christopher Chambers. 2012. Skill progression demonstrated by users in the Scratch animation environment. *International Journal of Human-Computer Interaction* 28, 6 (2012), 383–398.
- [49] Samantha Shorey, Benjamin Mako Hill, and Samuel Woolley. 2020. From hanging out to figuring it out: Socializing online as a pathway to computational thinking. *New Media & Society* (2020), 1461444820923674.
- [50] Nischal Shrestha, Titus Barik, and Chris Parnin. 2021. Remote, but Connected: How# TidyTuesday Provides an Online Community of Practice for Data Scientists. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–31.
- [51] Georg von Krogha, S. Spaeth, and K. Lakhani. 2003. Community, joining, and specialization in open source software innovation: a case study. *Social Science Research Network* (2003).
- [52] April Yi Wang, Zihan Wu, Christopher Brooks, and Steve Oney. 2020. Callisto: Capturing the "Why" by Connecting Conversations with Computational Narratives. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [53] Etienne Wenger, Richard Arnold McDermott, and William Snyder. 2002. *Cultivating communities of practice: A guide to managing knowledge*. Harvard business press.
- [54] Jeannette M Wing. 2008. Computational thinking and thinking about computing. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 366, 1881 (2008), 3717–3725.
- [55] Seungwon Yang, Carlotta Domeniconi, Matt Reville, Mack Sweeney, Ben U Gelman, Chris Beckley, and Aditya Johri. 2015. Uncovering trajectories of informal learning in large online communities of creators. In *Proceedings of the Second (2015) ACM Conference on Learning@ Scale*. 131–140.

A APPENDIX

CT concepts	Scratch Blocks
Loops	forever, foreverIf, repeat, repeatUntil
Parallelism	startHatTriggered, eventHatTriggered, keyHatTriggered, mouseHatTriggered
Events	eventHatTriggered, keyHatTriggered, mouseHatTriggered, bounceOffEdge, turnAwayFromEdge, touching, touchingColor, colorSees, mousePressed, keyPressed, isLoud, sensor, sensorPressed, distanceTo
Conditionals	waitUntil, foreverIf, if, ifElse, repeatUntil, bounceOffEdge, turnAwayFromEdge, touching, touchingColor, colorSees, mousePressed, keyPressed, isLoud, sensor, sensorPressed, lessThan, equalTo, greaterThan, and, or, not, listContains
Operators	lessThan, equalTo, greaterThan, and, or, not, add, subtract, multiply, divide, pickRandomFromTo, concatenateWith, letterOf, stringLength, mod, round, abs, sqrt, sin, cos, tan, asin, acos, atan, ln, log, e [^] , 10 [^]
Data	setVarTo, changeVarBy, showVariable, hideVariable, readVariable, addToList, deleteLineOfList, insertAtOfList, setLineOfListTo, contentsOfList, getLineOfList, lineCountOfList, listContains

Table 4. CT concepts mapping to blocks in the Scratch programming language, adopted from Dasgupta et al. [12]. We use this mapping to construct our outcome variable for H1a and H1b about whether new CT concept is presented in projects created after a user's first 14 days in the community.

Received July 2021; revised November 2021; accepted April 2022

Variables	Total new CT concepts (M1)	Active duration (M2)	New loves per project (M3)
(Intercept)	4.65** (0.01)	4.81** (0.01)	0.15** (0.00)
Made_original_project?	-0.09* (0.04)	-0.28** (0.02)	-0.15** (0.01)
Made_remix?	0.02 (0.02)	-0.04* (0.01)	-0.04** (0.00)
Made_friend?	-0.18** (0.02)	-0.03* (0.01)	0.02** (0.00)
Made_comment?	-0.09** (0.02)	0.32** (0.01)	0.16** (0.00)
Has_been_friended?	-0.06* (0.02)	0.09** (0.01)	0.08** (0.00)
Has_been_loved?	-0.03 (0.03)	0.11** (0.01)	0.19** (0.01)
Has_been_viewed?	0.00 (0.03)	0.06** (0.02)	0.00 (0.01)
Has_been_remixed?	-0.02 (0.03)	0.11** (0.02)	0.05** (0.01)
Has_been_commented?	0.01 (0.02)	0.14** (0.01)	0.01* (0.00)
Has_been_favorited?	0.01 (0.03)	0.11** (0.02)	0.12** (0.01)
log1p(CT_concepts)	-1.68** (0.02)	NA NA	0.03** (0.00)
R ²	0.47	NA	0.15
Adj. R ²	0.47	NA	0.15
AIC	NA	1140546.37	NA
BIC	NA	1140660.20	NA
Log Likelihood	NA	-570261.19	NA
Deviance	NA	116442.92	NA
McFadden's pseudo R-squared	0.16	0.01	0.14
Num. obs.	47952	97295	79963

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Table 5. Results from our early exploratory analysis based on continuous construct of outcome variables and subsets of our user level dataset used from the main analysis. M1_c tests H1, where *Total_new_CT_concepts* is a count variable of number of new CT concepts used in projects created after the first 14 days among the 47,952 users who had used new CT concepts after 14 days. M2_c tests H2, where *Active_duration* is a count variable of the number of days that a user engages in any recorded community activities and is measured by the number of days between the the end of the users 14 day period and the last in which they are active. This is also constructed on the subset of 97,295 users who were active after the first 14 days. M3_c tests H3, where *New_loves_per_project* is a continuous variable of average number of loves received by a user on projects created after the first 14 days, among the 79,963 users who did create projects after 14 days. We chose to exclude these results from our main findings due to small sample size (< 10% of our original dataset), and truncation issues.

Variables	New CT concept? (M1)	Stayed? (M2)	Received new love? (M3)
(Intercept)	1.83*** (0.02)	3.39*** (0.04)	-1.04*** (0.02)
Made_original_project?	-0.79*** (0.04)	-3.05*** (0.05)	-1.63*** (0.06)
Made_remix?	-2.34*** (0.04)	-3.59*** (0.05)	-1.54*** (0.06)
Made_friend?	-0.08 (0.05)	0.30*** (0.07)	0.21*** (0.04)
Made_comment?	0.03 (0.04)	0.52*** (0.06)	1.01*** (0.03)
Has_been_friended?	0.17*** (0.02)	0.35*** (0.02)	0.40*** (0.02)
Has_been_loved?	-0.05 (0.03)	0.16*** (0.02)	0.68*** (0.02)
Has_been_viewed?	-0.27*** (0.03)	-0.35*** (0.03)	0.14** (0.05)
Has_been_remixed?	-0.01 (0.03)	0.08* (0.03)	0.29*** (0.03)
Has_been_commented?	-0.13*** (0.02)	-0.01 (0.02)	0.25*** (0.02)
Has_been_favoriated?	-0.09** (0.03)	0.13*** (0.03)	0.39*** (0.02)
log1p(CT_concepts)	-1.22*** (0.02)	NA NA	0.13*** (0.02)
Made_comment? × Made_remix?	0.22*** (0.04)	0.17*** (0.04)	0.11** (0.04)
Made_comment? × Made_friend?	-0.01 (0.04)	0.11** (0.04)	-0.12*** (0.04)
Made_comment? × Made_original_project?	-0.20*** (0.04)	-0.19*** (0.06)	-0.10** (0.04)
Made_remix? × Made_friend?	0.10* (0.04)	0.05 (0.04)	0.15*** (0.04)
Made_remix? × Made_original_project?	2.30*** (0.05)	3.60*** (0.05)	1.59*** (0.06)
Made_friend? × Made_original_project?	0.06 (0.04)	-0.06 (0.06)	0.13*** (0.04)
AIC	122375.61	128298.77	118346.10
BIC	122550.30	128463.75	118520.78
Log Likelihood	-61169.80	-64132.39	-59155.05
Deviance	122339.61	128264.77	118310.10
McFadden's pseudo R-squared	0.25	0.16	0.14
Num. obs.	121149	121149	121149

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Table 6. Logistic regression models that include two-way interaction between independent variables as predictors. Same as our main analysis, the models predict the likelihood that a user using new CT concept (M1), staying active (M2), and receiving loves on new projects (M3) created after their first 14 days in the community. Models are fit on the user level dataset including aggregated activities from 121149 users.