"I Think It's Her": Investigating Smartphone Users' Speculation about Phone Notifications and Its Influence on Attendance

Yung-Ju Chang[1], Yi-Ju Chung[2], Yi-Hao Shih[1]

[1]National Chiao Tung University, Department of Computer Science, Hsinchu, Taiwan
{armuro, sam1424061933155.cs05g}@g2.nctu.edu.tw
[2]National Taiwan University, IoX Center, Taipei, Taiwan, violetachung@ntu.edu.tw

ABSTRACT

Smartphone users' decisions about whether to attend to a notification after sensing it are under-researched. We therefore studied 33 Android users, and found that they speculated extensively about notifications' sources—i.e., which apps and which senders were responsible for them—before attending to them. The participants' speculation about apps was both more common and more accurate than that about senders. They also were more likely to 1) perceive notifications as important, 2) attend to them, and 3) consider them beneficial if they speculated about them than if they did not or could not. Participants' speculations were based on the alert's inherent characteristics, context, and temporality. Inaccurate speculations were mainly caused by unclear signals, insufficient clues, and a multiplicity of possible sources. Ringer mode affected the accuracy of user speculation, but not its frequency or the frequency of attending to notifications.

CCS CONCEPTS

• Human-centered computing → Empirical studies in ubiquitous and mobile computing

KEYWORDS

Mobile notifications; mobile receptivity; speculation; ESM

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous;

MobileHCI '19, October 1–4, 2019, Taipei, Taiwan

ACM ISBN 978-1-4503-6825-4/19/10...$15.00
https://doi.org/10.1145/3338286.3340125

ACM Reference format:
Yung-Ju Chang, Yi-Ju Chung, Yi-Hao Shih. 2019. "I Think It's Her": Investigating Smartphone Users' Speculation about Phone Notifications and Its Influence on Attendance. In *Proceedings of the 21th International Conference on Human-Computer Interaction with Mobile Devices and Services* (MobileHCI, Oct 1–4, 2019, Taipei, Taiwan. ACM, New York, NY, USA.
https://doi.org/10.1145/3338286.3340125

1 INTRODUCTION

Notifications are a common means of receiving information from or via smartphones. However, users do not value all the notifications that they receive. Research has shown that notifications from certain apps or people are particularly valued and preferred by users [35,39,43]: at specific moments, users are likely to perceive notifications from those sources as beneficial to view, and other notifications as irrelevant and distracting. Unfortunately, as smartphone users install more apps on their phones, their attention is more likely to be drawn by alerts from notifications that they consider irrelevant. We assume that, when a user senses an alert—i.e., hears the sound or feels the vibration—he or she will often speculate about which app or sender is sending it and why, and base his/her decision about whether to attend to it on such thoughts, which we refer to as *notification speculation*. However, neither sound-nor vibration-based notification alerts are always a sufficient basis for accurate speculation. Although some mobile apps include custom alerts to help users distinguish them from other apps, many still use default tones and vibrations, or very similar ones. Despite the practical, commercial, and social importance of notification speculation, however, the literature on users' current practices and challenges in this area remains sparse. In building our understanding of this phenomenon and exploring the notification-alert design space with the wider aim of supporting more effective notification speculation, we will be guided by the following research questions:

RQ1: *How frequent and accurate are notification speculations, and what factors affect their frequency and accuracy?*
RQ2: *How do users speculate about notifications, and what are the challenges to making accurate speculations?*
RQ3: *How does the process of speculation affect users' decisions of and the effectiveness of their notification attendance?*

Our two-week empirical study utilized mixed methods including the experience sampling method (ESM), phone logging, and cued retrospective interviews. It makes four main contributions to the literature. First, it provides evidence of users' notification-speculation behavior and performance, and the effect of their speculation on subsequent attendance. Second, it identifies the major bases of users' speculations, along with situations that tend to make speculation difficult and/or inaccurate. Third, it reveals that ringer mode affects the accuracy of speculation about which apps are sending notifications, but not the frequency of such speculation or the frequency of attending to such notifications. Finally, it provides recommendations for the design of future notification systems that will help support notification speculation.

2 RELATED WORK

The body of mobile-notification research provides us with a rich understanding of smartphone users' notification-management and notification-attendance behaviors. Notification management comprises both which notifications mobile users prefer to see, and how they deal with them. People's receptivity to notifications has been found to be affected by how interesting and relevant the notifications are [9]. When smartphone users consider received notifications to be irrelevant, they often dismiss them and may even delete the app that sent them [35]. For example, they particularly value notifications related to people, such as those sent by from Instant Messenger (IM) apps [30,35]. When smartphone users chose to "snooze" notifications related to people and events, it has been taken to indicate a desire to view them later despite being currently unavailable [44]. Thus, the idea of taking users' preferences into account to determine what notification to deliver has been implemented by researchers [20].

Several studies investigated smartphone users' awareness and attentiveness to notifications. It is suggested that smartphone users are attentive to notifications, even without notification alerts [6,22]. For example, participants in [22] reported that they only missed notification alerts 14.63% of the time when their phone was silent. Other researchers used the notion of *attentiveness*, which quantified how quickly participants saw notifications and/or how frequently they attended to notifications in a given period. Both [6,7] found that, on average, participants attended to notifications within several minutes of the notifications arriving. The same studies' participants were also found to be attentive across all ringer modes, which on the Android system are of three types *Silent Mode* (the phone delivers no alert), *Vibrate Mode* (the phone delivers only a vibration alert), and *Normal Mode* (the phone delivers both sound and vibration alerts). Mashhdi et al. found that if a notification was associated with an alert, it was 12 times more likely to be attended to immediately than if it was not [19]. Other studies likewise reported that notifications that arrived in Silent Mode took the longest to be seen [22,27], whereas those that arrived in Vibrate Mode evoked the quickest responses [6] However, while users can feel more productive and less distracted when their alerts are turned off, they may actually feel worried about missing notifications [2,28,29], and thus self-initiate app-checking behaviors that render them more likely to delay or even forget their tasks on the phone [42].

Researchers have also attempted to identified factors correlated with or predictive of smartphone users' attentiveness and responsiveness to notifications, and to predict their interruptible or opportune moment for receiving notifications or performing phone-related tasks. These matters have been found to be influenced/predicted by or correlated with activity type and level [1,4,5,7,8,14], location [21,26,36,41], time of day [26,31,33,36,41], the recentness of users' interactions with their phones [8,14,16], ringer mode [9,14], sensor information [9,14-16,18,22], conversational context [37,38], personal traits [46], and arousal and emotional states of arousal [13,23]. In particular, [17,21,22] suggest that who the sender is also matters. However, if an intelligent system for helping users decide which notifications to attend/respond to is to be successfully developed, it will first be essential to elucidate *how* users speculate and make attendance decisions *prior to* attending to them, and therefore when such speculation may require assistance. [5] presents preliminary results in this regard, but it has limited quantitative and qualitative findings. This paper provides more such insights.

3 METHODOLOGY

3.1 Research Focus and Participant Recruitment

Based on its research focus as set forth above, this study's target population consisted of smartphone users who are usually informed of phone-related events via notification alerts, and its target scenario was users making an attendance decision after sensing a notification alert when the phone was not in use. According to prior research on mobile notifications, overall usage of non-Silent Modes is higher than the usage of Silent Mode on Android systems (e.g., [6,11,19,22,45]). Thus, this study's scope covers a considerable proportion of smartphone users. We selected the Android system as our research platform because it allowed us to capture behavioral data, such as app usage and ringer-mode settings, among other user/usage details.

We recruited participants who 1) set their phones to Silent Mode for not more than eight hours per day; 2) used more than one communication app, and at least two apps that used different notification alerts; and 3) were aware that their apps had different notification alerts. We balanced the participants in terms of gender, occupation, and self-reported number of daily notifications. We posted a recruitment message in a subject pool created at our own and a neighboring university; on a bulletin board intended for recruiting participants; and on the research team members' social-media pages. Initially, 37 participants were recruited, and numbered P1-P37. Data collection took place from March 17 to July 31, 2017. Three participants withdrew because the research app did not function properly on their phones. Thus, a total of 34 people (17 males, 17 females) completed the study, and of these, all but one participated in an optional interview. The participants were from various cities in Taiwan, mainly Taipei and Hsinchu, where people have a relatively high acceptance of technology. All were 20-36 years old.

3.2 Experience Sampling Study

We used ESM to study the participants' reactions to and experience of notifications. ESM prompts were triggered only after participants had started using their phones, so that the alerts from such prompts would not affect their sensing of and speculation about app notifications. We set a minimum interval of 90 minutes between any two ESM prompts. Whenever a participant saw an ESM prompt, the research app sampled between one and three notifications that had arrived in the 30 minutes before the ESM prompt. Based on feedback from our internal pilot test, we determined that this 30-minute threshold was long enough to capture diverse sensing and speculating scenarios, but not too long for clear recall. It was crucial to include not only notifications arriving right before the ESM prompt (i.e. when participants just attended to the phone), because

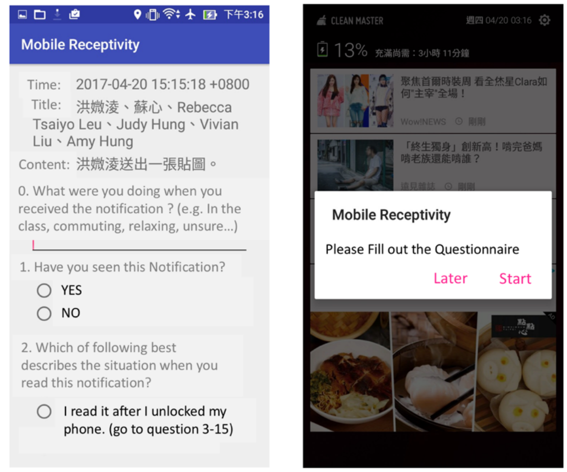

Figure 1. (Right) The ESM prompt. (Left) The ESM questionnaire, containing the title, content and time the notification was received, followed by a set of questions inquiring about the participant's experiences.

doing so would have prevented us from capturing situations in which the participant decided not to attend to the phone after speculation. Given the target research scenario, the research app did not prompt participants about notifications that arrived either when they were already using the phone for some other purpose, or when the phone was in Silent Mode. Each selected notification had its own questionnaire (see Figure 1), which contained its title, content, and arrival time, followed by 16 multiple-choice questions covering sensing; speculation (action, basis for speculation, and correctness); reasons for attending; and self-evaluation of attendance. For participants unsure if they had sensed a notification, an "unsure" option was provided. Due to space limitations, the questionnaire is presented in English translation in the supplemental material, rather than the original Mandarin.

3.3 Study Procedure and Data Collection

Participants in the final sample ran our research app on their phones for 14 days. It logged notifications, the participants' actions, and phone-context information including location, activity, phone sensor, and phone status (e.g., network connectivity, battery life). We used the Android system's Accessibility Service[1] to record participants' actions and infer their attendance to notifications, in line with prior research (e.g., [6,7,22]). In a pre-study meeting, we sought the participants' informed consent, and if they agreed to participate, helped them to install the research app and provided them with a tutorial on answering ESM questionnaires. We asked the

[1]https://developer.android.com/reference/android/accessibilityservice/Accessibility Service.html

participants to connect their phones to wi-fi daily to upload logged data. Upon completion of the ESM study, we provided participants with a gratuity of 1,200 NTD (about 40 USD) and invited them to an interview via email, with attachments containing statistics of their notifications and ESM responses. We asked them to review this data and to recall its contexts prior to the interview. In the interview sessions, we used a cued retrospective technique [34], which involved showing the interviewee his/her notifications and responses, to obtain contextual details and explanations of these instances. All interviews were audio-recorded and transcribed. The interview was optional; each participant who took part in an interview received an gratuity of 300 NTD (10 USD).

3.4 Data Cleaning and Analysis

We collected 713,866 notification events in total, but a large proportion of these were "ongoing notifications" that did not generate alerts (e.g., navigation status; music playing), and were thus removed from data analysis. In keeping with prior research that quantified attentiveness (e.g., [6,7,22]), we measured notification-attendance time as the time elapsed between notification arrival and the next logged action or next "Interactive" status of the phone, whichever came first. The actions of dismissing and clicking were both included, as both require users to switch attention. Because prior research has indicated that smartphone users, on average, returned to an attentive state within 5 minutes [7], we applied a threshold of 5 minutes to differentiate between attended-to and non-attended-to notifications. We also included the 1-minute threshold used in one previous study [6] as an alternate measure of immediate attending.

For ESM, we removed P35's ESM responses because, in the interview, we found that he critically misunderstood several ESM questions. In the rest of 4,412 responses, we removed ESM responses to notifications not in the target scenario (i.e. received when the phone was in Silent Mode and when it was screen-on), resulting in the final 4212 ESM responses. We inspected the validity and the logical consistency of ESM responses. Because the validity of responses to many questions depended on the participants' responses to earlier questions (e.g. answering speculation related questions only after reporting having sensed the notification earlier), for each question we set filters on its earlier questions in a spreadsheet so that we could filter out invalid responses to that question.

When examining whether an observed pattern was statistically meaningful, we employed mixed-effects regression analysis because each participant had an imbalanced number of ESM data. A logistic regression model was chosen when the dependent variable (DV) was a binary variable (e.g. speculated versus not speculated), and a linear regression model when it was a continuous variable (e.g. notification seen time). User, hour of the day, and day of the week were included as random effects. For app types, we used the categories devised by Shirazi et al. [35]. When analyzing the effect of app type included as an independent variable (IV) in the regression model, we considered only those app types for which we had more than 20 observations. For qualitative data, we conducted a thematic analysis [4]. Two co-authors independently transcribed the interview recordings and coded the transcripts using an iterative process of generating, refining, and probing emergent themes. A third author regularly met with the two coders to discuss the codebook and the high-level themes. The coders iteratively revised the codebook until, in a particular iteration, they reached an inter-rater reliability of 0.86 using Cohen's Kappa [16]. They then used the codebook to code the rest of the data and identify its high-level themes.

4 QUANTITATIVE RESULTS

We analyzed 168,262 notifications in total. The number of notifications received by individual participants varied from 1,600 to 11,243 (M=4,949, SD=2,711). The top five categories of notifications were IM (51.7%), system (17.1%), utility (14.4%), mail (5.2%), and social (4.7%). 38.3% of all notifications were collected in Normal Mode, 57.1% in Vibrate Mode, and 4.6% in Silent Mode. Given our participant-recruitment strategy, the small proportion of notifications received in Silent Mode was expected. Among the 4,212 ESM responses we analyzed, 65.8% reported that the user had seen the notification before he/she filled out the questionnaire. Of these seen notifications, 71.1% (1,869) were seen after participants attended to their phone; and of these 1,869 responses, 52.4% (931) were reported as sensed. It was this group of 931 sensed notifications that reflected the target scenario. The numbers of the apps that generated these notifications per app category were 11 (Messenger), 10 (Social), 6 (mail), 6 (reader/news), 6 (systems), 13 (utility) among the 931 instances.

4.1 Speculation is Common and Generally Accurate

The participants speculated in 71.6% (667) of the 931 notification instances that matched our target scenario. Among the 667 instances, the speculation category included both thoughts that a given notification was from a specific app and/or sender (93.1%, 621 out of 667), and a reported inability to guess what its source was (6.9%, 46 out of 667). In particular, the participants speculated *both* about

Table 1. We group eleven phone-event types into three categories – User Presence, Receiving Notification, and App Use.

Type	Measure	Detected Event Content
User Presence	User Interaction	Existing user action
	Screen On	Existing screen-on event
Receiving Notification	Recv Noti	The phone received any notifications
	Recv Msg Noti	The phone received message notifications
	Recv Same Cat App Noti	The phone received notifications in the same category
	Recv Same App Noti	The phone received notifications from the same app
App Use	Used App Same As Noti	The participant had used the app that was the notification source
	Used App Cat Same As Noti	The participant had used an app whose category was the same as the notification source
	Used Msg App	The participant had used a messaging app
	Used Msg App Cat Same As Noti	The participant had used the same messaging app that produced the notification
	Used Msg App Same As Noti	The participant had used a messaging app whose category was the same as the notification source

the apps and the senders 49.6% (331 out of 667) of the time, as compared to 39.7% (265 out of 667) of the time for the apps only, and 3.7% (25 out of 667) of the time for the senders only. These figures indicate the participants speculated more frequently about apps (89.4%, 596 out of 667) than about senders (53.4%, 356 out of 667). Alarm clock apps were always speculated about. The next most frequently speculated-about app types were IM (72.26%). Using app as an IV in a logistic regression model with IM as a reference class, Wald tests of each app type show that IM apps were more often speculated than social (54.6%, z =-4.2 , p<.001), utility (48.8%, z=-2.2 , p=.02), reader/news (46.5%, z=-3.8 , p<.001), mail (43.4%, z=-5.2, p<.001), and system (40.9%, z=-2.7, p=.006).

4.1.1 *App Speculation More Accurate than Sender Speculation.* When the participants were able to speculate about their notifications, i.e., those for which they felt they could tell the source, their speculations were highly accurate: 89.5% on average. However, their speculations about apps were more accurate than those about senders (96.2% vs. 86.2%). Among the 330 instances in which participants speculated about both, they were incorrect about the sender 34 times, but about the app only once. In addition, when participants speculated only about the sender, their accuracy was only 45.5%. Although only 11 speculations about just the sender were made, the data appears to indicate that accurate speculation about senders was more challenging than accurate speculation about apps.

4.1.2 *Speculation about IM Apps the Most Accurate.* The participants' speculations about apps were more accurate for IM notifications (98.8%) than for social (87.3%, z=-3.7, p<.001) and mail (84.4%, z=-3.2, p<.001) notifications. When speculating about the sender of a notification, the participants' accuracy was lowest for mail apps (53.3%). The differences between mail apps and social, and between mail apps and IM apps, respectively, were both statistically significant (vs. social: 96.3%, z=2.4, p=.02; vs. IM: 87.3%, z=4.0, p<.001). This was probably because senders of e-mails were both more diverse and less likely to engage in continuous interaction.

4.2 Speculation is Related to Recent Phone Events

We examined whether either the likelihood or the accuracy of participants' speculations was related to a *phone event* that happened prior to the arrival of the notification. As shown in Table 1, we examined 11 phone-event types, grouped into three broader categories – *User Presence* (using the phone, screen-on); *Receiving Notification*; and *App Use* – and six time windows: *<1 min, <5 min, <10 min, <30 min, <60 min,* and *>60 min.* For example, *Used App Same as Noti <10 min* implies that the participant had used the app that generated the target notification within 10 minutes prior to receiving that notification. Figure 2 shows the results of the eight most distinct phone events, which also includes an average of the measure (regardless of the existence of the phone event and the time), plotted in red as a reference line. The figure does not include app-speculation accuracy because, as noted above, this was very high overall. As the plot shows, recent phone events were associated with a higher likelihood 1) that speculation would occur, and 2) that it would be correct; and the more recent these events were, the more true this was. In addition, specific categories of phone events were especially indicative of certain levels of likelihood. For

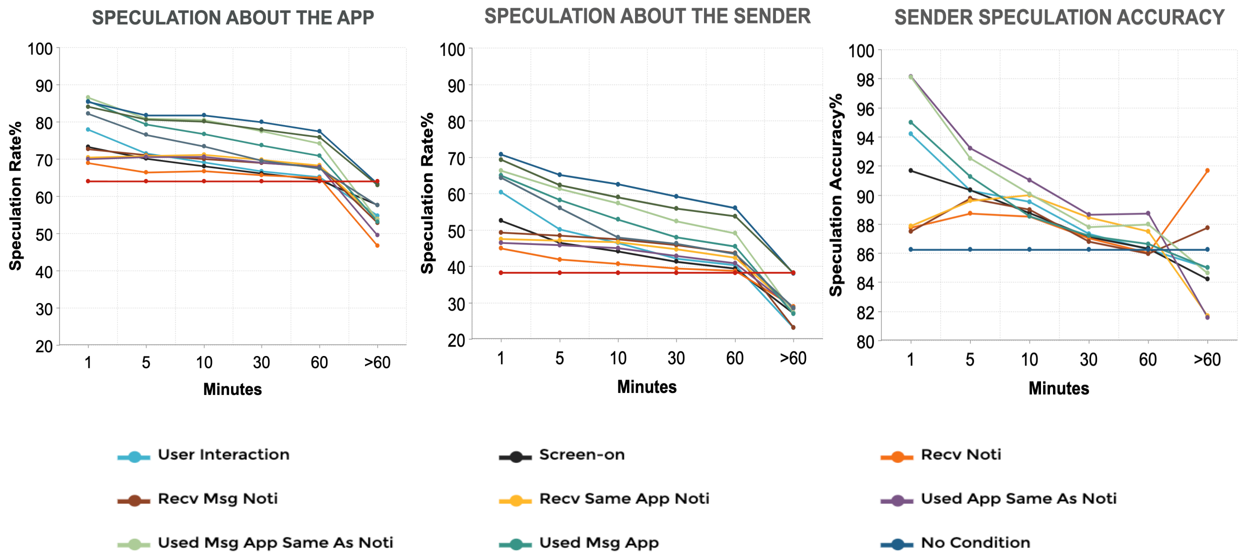

Figure 2. Three phone-event types—Recent User Presence, Recent Receiving Noti, and Recent App Use—that occurred within one hour prior to the arrival of the target notification, as indicators of the participants' speculation about the app (left), speculation about the sender (center), and sender-speculation accuracy (right).

example, recent *App Use* was associated with a high likelihood that a notification would be speculated about, and correctly speculated about; and *Used Msg App Same as Noti* was especially prominent in this category. Such results imply that, although the participants might have checked their phones when starting new tasks [3], their recent use of the same app that generated that notification rendered them more likely to speculate, and to speculate correctly. Conversely, the absence of *App Use* events seemed to be less indicative of a low likelihood of speculation than the absence of *Receiving Noti* and *User Presence* events were. This implies that when participants had not received a notification or had not interacted with the phone for more than an hour, there was a low likelihood that they speculated about the notification source when receiving a notification. It is noteworthy that, however, *Receiving Noti* was the least clearly linked to the occurrence of speculation behavior than any of the phone-event types, possibly because participants were used to notifications constantly arriving. As these results suggest an association between notification speculation and recent phone events, later qualitative data also provide support of such an association.

4.3 Speculation's Influence on Attendance

We analyzed speculation's influence on both logged attendance measure (LM) and ESM self-reported attendance measure (EM). Although these two measures differed fundamentally (one being based on inference, and the other, subjective), both suggested that the participants' speculation affected their subsequent attendance.

4.3.1 Speculation and the Likelihood of Attending. As Figure 3 shows, the participants' overall attentiveness to notifications was higher when they could speculate about them, regardless of correctly or not, (EM: 80.0%; LM 5-min: 83.4%; LM 1-min: 67.7%) than otherwise, including when they were not able to tell the source (EM: 78.3%; LM 5-min: 68.9%; LM 1-min: 48.9%) and when they did not speculate (EM: 63.6%; LM 5-min: 71.7%; LM 1-min: 49.6%). The differences for all three attendance measures were statistically significant (EM: z=5.0, p<0.001; LM: z=3.7; LM 1-minute: z =5.1, p<0.001). This indicates that when participants could think of a source before attending to the notifications, they were more likely to attend to them than when they did not and could not think of a source,

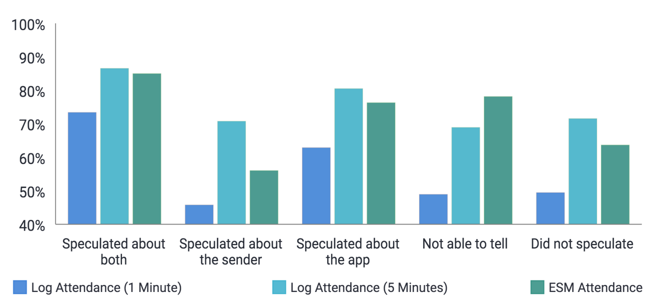

Figure 3. Notifications being attended to according to participants' speculations.

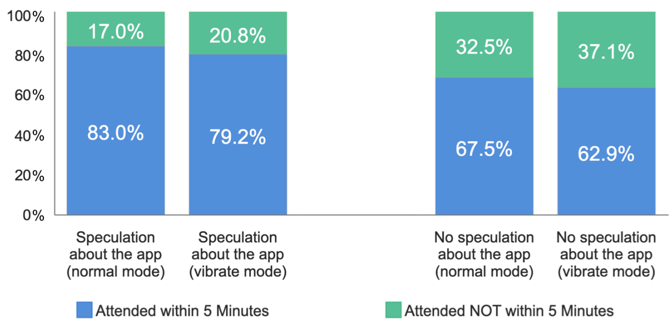

Figure 4. Percentage of notifications attended to within 5 minutes, by app speculations and ringer modes.

respectively. On the other hand, the participants seemed to be most attentive when they speculated about both the app and the sender (EM: 84.9%; LM 5-min: 86.7%; LM 1-min: 73.3%). In addition, when their speculations involved an app (EM: 81.0%; LM 5-min: 83.9%; LM 1-min: 68.7%), they were more attentive than when the speculations did not involve an app (EM: 56.0%; LM 5-min: 70%; LM 1-min: 45.8%). These differences were also all statistically significant (EM: z=5.0, $p<0.001$; LM: z=3.6; LM 1-minute: z=5.0, $p<0.001$). This finding suggests that whether the participants could think of the app of the notification played a critical role in their subsequent decision of whether to attend to the notification.

4.3.2 Attentiveness by Ringer Mode and App Type.

A number of prior studies have linked alert modality to users' attentiveness to notifications [6,22,27]. However, our results do not show an effect of ringer mode in our target scenario. Specifically, when considering all notifications, including those received not in the target scenario, notifications received in Vibrate Mode were seen faster than those received in Normal Mode ($t(149,069)=-8.4$, $p<0.001$). However, when considering only notifications in the target scenario, there were no statistically significant differences in attentiveness between Normal Mode and Vibrate Mode in either EM (z=0.9, p=.35), LM (5-min: z=1.1, p=.28; 1-min: z=1.1, p=.26), or logged attendance time ($t(428)=0.66$, p=.51) when speculation was taken into account. As shown in Figure. 4, attentiveness to notifications was mainly affected by whether or not participants had speculated about the app, rather than by ringer mode. Moreover, there was only a slight difference in attentiveness across ringer modes, regardless of speculation. The fact that the participants did not speculate more often in one ringer mode than another explains why they also did not attend more in a particular mode. However, the differences between speculation about an app occurring and not were larger than 15% for both of the non-silent ringer modes. This shows that what mattered to attentiveness in the target scenario was not the ringer mode,

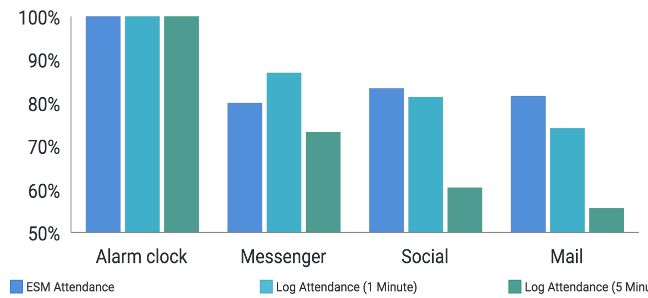

Figure 5. Notifications attended to, by app type.

but whether the participant speculated about it. On the other hand, we observed disparities in the participants' attentiveness to notifications from different app types across EM and LM. Among correctly speculated notifications (see Figure 5), attentiveness measured by LM and EM noticeably differed. For example, LM data appeared that the participants were more attentive to IM apps (LM 5-min: 86.9%, LM 1-min 73.1%) than to both social apps (LM 5-min: 81.3%, LM 1-min 60.4%) and mail apps (LM 5-min: 74.1%, LM 1-min 55.6%). However, this pattern was absent in EM (IM: 79.9%, social: 83.3%, mail: 81.5%). Unfortunately, we could not conclude whether the disparities were because of the different nature between EM and LM or because of the lack of effect of app type on attentiveness.

4.3.3 Speculation and Perceived Importance of Notifications

Finally, in ESM, participants answered multiple-choice items on their reasons for attending and not attending to notifications, regardless of whether they *speculated* (successfully thinking of a specific source), *did not speculate*, or *could not tell the source* (attempting to speculate but fail to think of a specific source). Note that here the outcome of speculations is not considered, since we are mainly interested in participants' attendance decision after reacting to notification arrival. *I had time to read it* was the top reason for attending across all speculation categories (speculated: 64.6%, did not speculate: 68.4%, not able to tell: 69.4%). *I did not have time to read it* was the top reason for not attending for speculated (57.7%) and did not speculate (52.1%). However, when participants could not tell the source, they chose *the notification was unimportant* 50% of the time. While this sample was small, it is interesting that *the notification was unimportant* was chosen only 9.8% of the time by those who speculated, and just 13.8% of the time by those who did not speculate. On the other hand, *the notification was important* was chosen as the third most important reason for attending to among those who speculated (24.8%). The same reason was chosen only 12% and 13.9% of the time by those who did not speculate (12.0%) and those who could not tell the source (13.9%), respectively. This implies that when participants speculated about its source, they were likely to regard a notification as important, and thus to attend to it. Taken together, these findings provide some explanation about why participants' overall attentiveness to notifications was higher when they speculated about them than otherwise.

4.3.4 Speculation and the Effectiveness of Attendance. We asked the participants to rate whether their attendance or non-attendance to each notification was beneficial to them at the time, using one of the following three options: 1) Necessary; 2) Beneficial to Read; and 3) Unnecessary. When participants had speculated about a given notification's source, they considered their attendance beneficial 90% of the time (i.e., Necessary: 48.2%, Beneficial to Read: 41.8%). In contrast, when they could not tell the source or did not speculate, they rated the attendance as Unnecessary 51.9%; and 35.7% of the time, respectively. Interestingly, when the participants decided to not read notifications after sensing them, they rarely considered those unattended notifications beneficial to read: with only five out of all 212 such instances (2.4%) being assigned this rating. This finding shows an interesting contrast to the concern of missing important notifications found in [2,28,29] when the phone is silent. Perhaps when the phone delivers an alert, users would feel having more control of what to read and what not to read and thus would not think that their decision of not reading a notification had caused them to miss something important.

5 QUALITATIVE FINDINGS

5.1 Bases of Speculation: Alert, Context, Temporality

Our analysis of mixed data sources revealed that the participants' speculations had three main bases: *notification alert, context,* and *perceived temporality.* In their ESM responses, the participants were asked to choose all options on which their speculations about the app were based, but regarding their speculations about senders, they chose only the major basis. The difference in the provided options was based on feedback we had obtained from our pilot study.

5.1.1 Notification Alert

Given that many apps incorporate their own custom ringtones or vibration alerts, and that we actively sought to recruit participants who were aware of differences between notification alerts, it was unsurprising that many of them said they commonly distinguished between apps based on their alert patterns. As P29 stated, apps *"have different vibrations. You can clearly tell which app it is."* Moreover, after speculating about the app, the participants could sometimes also infer the sender of the notification and/or message content, based on their perceptions of what such content typically is: e.g., casual chat from a certain IM app, or promotional news from a commercial one. For example, P9 commented: "*I'd especially take a look at messages from Slack. [...] When I heard it I'd take a look, because the messages there are usually more important.*" The participants' tendency to associate certain people or meanings with a particular IM app was also aligned with *communication place* discussed in [24].

Additionally, the ESM and phone log data seemed to display three interesting associations between speculation and ringer mode. First, the participants reported a similar frequency of notification speculation across Normal Mode and Vibrate Mode (71.1% vs. 72.8%). Second, even when a sound alert was present (i.e., in Normal Mode), the participants often reported speculating about the app based on the vibration pattern. Specifically, among participants whose phones were in Vibrate Mode, *special vibration* was chosen as the basis of speculation in the ESM 66.7% of the time, whereas *special sound* was chosen only 11.2% of the time. Some participants chose *special sound* because they were referring to the sound generated by vibration. In Normal Mode, in contrast, *special sound* was chosen only 59.1% of the time, and *special vibration* as much as 44.3% of the time. This result implies the participants' reliance on haptic feedback even when sound alerts were enabled. Conceivably, this reflected their use of phones in noisy situations where ringtone-based alerts were not as noticeable as vibration-based ones. However, it is important to note that Android OS allows users to disable notification alert modalities for individual applications or within an application setting [15]. It was likely that some participants chose vibration as the speculation basis in Normal Mode because the sound of that notification had been suppressed. Since it is unclear how many such responses existed in our dataset, the observed percentage may not be conclusive. Finally, the participants' app-related speculations were more accurate when their phones were in Normal Mode (98.9%) than when they were in Vibrate Mode (93.5%) (p=.02). While this seemed to suggest that sound alerts were more, not less, distinguishable than vibration alerts, it could also simply be because Normal Mode provides more information than Vibrate Mode does.

5.1.2 Context

While notification alerts were mainly linked to apps, the participants often speculated about the sender or aboutness of the arriving notifications based on the context in which or about which they expected to receive a notification. They primarily speculated based on two types of context: a *recent phone event*, and a *perceived situation*. Within the former category, they often selected recent interaction as a basis for speculation about the app (41.1%; ranked 2nd) and about the sender (79.5%; ranked 1st). In the interviews, all but two participants mentioned this basis. As P4 stated: *"He had sent several messages earlier. I continued to watch TV. Then, I heard the sound [… and] thought it was him."* In addition, participants also speculated based on the *perceived situation* in which, or about which, they expected to receive a notification. For example, P2 reported: *"We're going to meet up later; he knew when to send me a message and ask where I am."* Similarly, P21 anticipated a notification after making a payment: *"So, this was when I paid with my card, I heard its 'dingdong'. Then, I guessed it was that app. Then, I opened it to see if the payment amount was right."* Other examples of successful app-related speculation involved a low-battery warning system and an alert indicating the detection of a new wi-fi signal.

5.1.3 Temporality: Frequency and Social-temporal Norms

Most participants also used *temporality*, especially *frequency* and *temporal pattern*, to speculate about notifications' sources. In other words, they tended to perceive both the people and the apps that sent them notifications as doing so at particular times and in particular quantities; and such perceptions were especially strong when such interactions were commonplace. Regarding temporal pattern, participants frequently mentioned that a social-temporal norm prevailed between them and their interlocutors. For example, P27 stated, *"If I didn't say anything about my dinner plans to my Mom in the morning, she would Line me and ask if I'm going home for dinner around time."* Similarly, P3 commented, *"my Mom typically looks for me around eight to do a webcam call. So, it's like a habit."* These communicational expectations, which can be described as micro-coordination [18], facilitated our participants' effective speculations about notification-senders' identities. Yet, temporal patterns also applied to speculation about apps. As P1 reported, *"the ones that ask you to download stickers. [… arrive] around eleven-fifty till twelve ten. They would be like, din-din-din-din-din. I'd just ignore them in that period."*

5.2 When Are Speculations Difficult or Incorrect?

While participants were generally accurate in their speculations, all of them sometimes made incorrect speculations or positively felt they had no way of guessing a particular notification's source. Both of these reactions were typically ascribed to *unclear signal*, *insufficient clues*, and *multiple possible sources*. With regard to the first, perceived clarity was affected by numerous factors, including the physical placement of the phone, environmental sound, and the participant's attention level. For example, P2 commented, *"The phone's vibration was absorbed by the bed. […] If you put it on something harder, like a table, it would be easier to tell."* The second category refers to situations in which the participant either could not associate the time he/she received the notification with any particular source or topic, or when the notification did not match his/her speculations. As P2 explained, *"my family would text me after I'm off work. I can guess that. Or when I just had a conversation with someone, I can guess that. But, friends? We're casual and random with texting. You can't be accurate."* As an example of receiving unexpected notifications, P15 stated, *"I bought something online with my credit card, and I expected to get a confirmation message from my bank. […] I heard the sound and thought that was the confirmation message. And then it turned out to be spam."* Lastly, a multiplicity of plausible notification sources also made speculation challenging: e.g., when the participant had been interacting with several people at a time, or associated the same time of day with different senders. As P27 put it, *"When I'm chatting with many people, and my phone only vibrates, I'm very likely to make a wrong guess. […] [S]ometimes, I would think that it's five in the afternoon and my Mom usually sends me a message. And so I'd just think it should be my Mom. But, it turned out to be my friend."*

6 DISCUSSION

6.1 Speculation is Vital to Notification Attendance

Our results indicate that notification speculation was prevalent in situations where the participants relied on notification alerts to stay informed about phone events when not using their phones. As well as being prevalent, however, notification speculation was found to influence individuals' decisions about when and whether to attend to notifications. Our participants were more likely to consider those notifications whose sources they had successfully speculated about as important than either those they failed to speculate about or those they did not speculate about at all. We also observed that, even when a participant had guessed only roughly where a notification was from and what it might be about, he/she knew whether it was necessary or beneficial to read at that moment. In contrast, when such guesses were off-base, the participants were much more likely to attend to notifications that they subsequently rated as unnecessary to read. With regard to the effect of app type, prior research has indicated that smartphone users prefer and are more attentive to notifications from certain apps [35,39,43], and our results indicate that such preferences were manifested in post-speculation attendance decisions. Among our participants, notifications from IM apps were more often successfully speculated about, attended to, and considered beneficial to read than those from any other app type. These results represent important evidence for speculation's key role in rendering notification attendance effective, at least in situations where users rely on alerts.

6.2 Notification Alerts' Impact on Attentiveness

Prior studies reported that notifications were seen fastest when phones were in Vibrate Mode [6,22,27]. However, we found no effect of ringer mode on attentiveness. While this result seems to conflict with prior findings, it may in fact provide explanations for ringer-mode effects that have not previously been articulated. As we noted earlier, when all notifications in our dataset were considered, those received in Vibrate Mode were seen faster than those received in Normal Mode. However, the full dataset included notifications that arrived when the phone was screen-on, and had very low seen-time. In our studies, these screen-on notifications were more frequently associated with Vibrate Mode (50.2%, i.e., 48,283 out of 96,152) than with Normal Mode (41.5%, i.e., 26,733 out of 64,368). Consequently, when these quickly attended-to notifications were eliminated from consideration as being irrelevant to the target scenario, such exclusion led to a lack of significant difference between ringer modes. As such, it would appear that previous scholars' observations of the effects of ringer mode on attentiveness might be attributable to the contexts in which their participants preferred using a quiet ringer mode (e.g., continually using the phone for chatting) rather than to the ringer *per se.* However, as mentioned earlier, another reason contributing to this inconsistency might be the difference between ringer mode and actual modality participants observed. Thus, the actual effect of alert on attentiveness would need more research to clarify. On the other hand, we have uncovered additional connections among ringer mode, speculation, and attentiveness: notably, that participants' attentiveness was driven by whether or not they had successfully speculated about which app the notification had come from, which in turn was related to alert clarity the participants perceived. Moreover, vibration remained a vital component of speculation even when a ringtone-based alert was also present (though the latter type of alert did increase the accuracy of the participants' speculations). We hoped that future alert systems take these findings into consideration.

6.3 Recommendations for Supporting Speculation

Based on the results, we feel that apps should be the first port of call in any effort to boost the effectiveness of user notification speculation. Although our participants' app-related notification speculations were quite accurate, this accuracy could have been overestimated due to hindsight bias. In addition, apps play a key role in non-app-related speculation: with participants' speculations about senders only being considerably more accurate when they involved consideration of the apps. We therefore propose the augmentation of notification alerts—for example, by extending their duration, amplifying them, or adding additional signal—when the system detects a situation in which users may have difficulty distinguishing among apps based on their normal alerts. However, it would be insufficient to only support speculation about apps, because users must also rely on context (recent interaction and perceived situation) and temporality (frequency and temporal pattern) if they are to speculate effectively about senders and notification content. However, these bases of speculation are not always available. Speculations can also be incorrect because notifications come from unexpected senders or because multiple possible associations exist. Therefore, to further assist users to form accurate and detailed previews of their notifications, we propose that future notification systems offer an additional signal when their systems detect that the current moment might impede correct speculation (e.g., situations in which the user could reasonably associate the same alert with various potential senders, or in which he/she receives a notification from an unfamiliar or otherwise unexpected source). Such solutions should follow guidelines for designing Earcons [12,32] that are easy to learn to reduce the burden of learning.

6.4 Study Limitations

It is important to recognize the limitations of the present study if its results are to be interpreted appropriately. First, it was based on a small sample of participants in Taiwan. The two most popular IM apps in Taiwan, Line and Facebook Messenger, may have different notification alerts from the IM apps that are most popular in various other countries. Consequently, it is unclear how generalizable our results are to other countries. In addition, our recruitment of participants could have been biased toward those who were most able to distinguish among various alerts and/or who were already accustomed to speculating about them. It is thus unclear how accurate the app-related notification speculations of the general population of smartphone users would be. Second, the ESM items covering the participants' attendance behavior used the term "right away," the meaning of which is subjective and could have varied sharply from one participant to another. Logged attendance, meanwhile, has limitations of its own. While we observed both consistencies and disparities between these two measures, it remains unclear which was closest to the ground truth. Third, showing participants notifications and then asking about their sensing of and speculation about them could potentially lead to hindsight bias [8,14]. Thus, we did not emphasize speculation accuracy, but attempted to discern overall trends and make comparisons between speculation types. To avoid this bias, future research may consider instructing participants to answer speculation questions before reading the notification. Fourth, to reduce recall bias in the ESM, we only sampled notifications that had been received within 30 minutes prior to the participant's most recent use of his/her phone. Given that the ESM questionnaire was lengthy, self-reported attentiveness could have been biased toward situations in which the participants were more attentive to their phones. Fifth, due to the specifications of our target scenario, we did not consider users who mainly kept their phones in Silent Mode. These users might display different kinds of speculation when they switch their phone to Normal or Vibrate Mode. Nor did this research consider alerts other than those based on ringtone and vibration, e.g., the flashing lights offered by certain phones. In addition, it seems likely that in certain environmental situations, the participants would have heard sounds produced by vibration in Vibrate Mode. As our research instruments did not capture the sounds that the participants actually heard, we cannot know the ground truth of what they sensed. Finally, a recent study showed that relationship affects users' attentiveness to IM notifications [17]. Speculating about who the sender was might affect participants' attendance decision, but this was not explored in this study.

7 CONCLUSION

Alert systems are intended to inform users of the arrival of notifications on their phones, but users do not always attend to notifications optimally. This paper's investigation of the prevalence and role of notification speculation among Taiwanese smartphone users who relied on notification alerts found that notification speculation was prevalent in our target scenario, and that it influenced 1) the participants' decisions to attend to notifications; 2) their actual attentiveness to notifications, and 3) the effectiveness of their attendance to notifications. We also found that ringer mode did not affect the frequency of either speculating or attending. However, it did affect the accuracy of speculation. Moreover, we determined that participants mainly used alert signal, context, and temporality to form their speculations; and inaccurate speculations were mainly attributable to unclear signal, insufficient clues about the source, and multiple possible notification sources. In sum, notification speculation is worth supporting in future notification systems, due to its prevalence and its influence upon notification attendance; and it is hoped that this paper's design recommendations will assist smartphone users to speculate about notifications more accurately, and to make better decisions about whether to attend to them or not.

ACKNOWLEDGMENTS

This research was supported in part by the Ministry of Science and Technology, R.O.C (MOST 107-2218-E-009-030-MY3, MOST 107-2218-E-009-053).

REFERENCES

[1] S. Aminikhanghahi, R. Fallahzadeh, M. Sawyer, D. J. Cook, and L. B. Holder. 2017. Thyme: Improving Smartphone Prompt Timing Through Activity Awareness. In 2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA), 315–322. DOI: https://doi.org/10.1109/ICMLA.2017.0-141

[2] Julie H. Aranda and Safia Baig. 2018. Toward "JOMO": The Joy of Missing out and the Freedom of Disconnecting. In Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '18), 19:1–19:8. DOI: https://doi.org/10.1145/3229434.3229468

[3] Niels van Berkel, Chu Luo, Theodoros Anagnostopoulos, Denzil Ferreira, Jorge Goncalves, Simo Hosio, and Vassilis Kostakos. 2016. A Systematic Assessment of Smartphone Usage Gaps. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16), 4711–4721. DOI:https://doi.org/10.1145/2858036.2858348

[4] Richard E. Boyatzis. 1998. Transforming Qualitative Information: Thematic Analysis and Code Development. SAGE.

[5] Yung-Ju Chang, Yi-Ju Chung, Yi-Hao Shih, Hsiu-Chi Chang, and Tzu-Hao Lin. 2017. What Do Smartphone Users Do when They Sense Phone Notifications? In Proceedings of the 2017 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2017 ACM International Symposium on Wearable Computers (UbiComp '17), 904–909. DOI: https://doi.org/10.1145/3123024.3124557

[6] Yung-Ju Chang and John C. Tang. 2015. Investigating Mobile Users' Ringer Mode Usage and Attentiveness and Responsiveness to Communication. In Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '15), 6–15. DOI: https://doi.org/10.1145/2785830.2785852

[7] Tilman Dingler and Martin Pielot. 2015. I'll be there for you: Quantifying Attentiveness towards Mobile Messaging. In Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services, 1–5.

[8] Florian Fessel, Kai Epstude, and Neal J. Roese. 2009. Hindsight bias redefined: It's about time. Organizational Behavior and Human Decision Processes 110, 1 (September 2009), 56–64. DOI: https://doi.org/10.1016/j.obhdp.2009.07.001

[9] Joel E. Fischer, Nick Yee, Victoria Bellotti, Nathan Good, Steve Benford, and Chris Greenhalgh. 2010. Effects of Content and Time of Delivery on Receptivity to Mobile Interruptions. In Proceedings of the 12th International Conference on Human Computer Interaction with Mobile Devices and Services (MobileHCI '10), 103–112. DOI: https://doi.org/10.1145/1851600.1851620

[10] Claudio Forlivesi and Utku Günay Acer. Mindful Interruptions: A Lightweight System for Managing Interruptibility on Wearables. 6.

[11] Jose A. Gallud and Ricardo Tesoriero. 2015. Smartphone Notifications: A Study on the Sound to Soundless Tendency. In Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct (MobileHCI '15), 819–824. DOI: https://doi.org/10.1145/2786567.2793706

[12] Stavros Garzonis, Simon Jones, Tim Jay, and Eamonn O'Neill. 2009. Auditory Icon and Earcon Mobile Service Notifications: Intuitiveness, Learnability, Memorability and Preference. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09), 1513–1522. DOI: https://doi.org/10.1145/1518701.1518932

[13] Nitesh Goyal and Susan R. Fussell. 2017. Intelligent Interruption Management Using Electro Dermal Activity Based Physiological Sensor for Collaborative Sensemaking. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. 1, 3 (September 2017), 52:1–52:21. DOI: https://doi.org/10.1145/3130917

[14] Scott A. Hawkins and Reid Hastie. 1990. Hindsight: Biased judgments of past events after the outcomes are known. Psychological bulletin 107, 3 (1990), 311.

[15] Andreas Komninos, Elton Frengkou, and John Garofalakis. 2018. Predicting User Responsiveness to Smartphone Notifications for Edge Computing. In Ambient Intelligence (Lecture Notes in Computer Science), 3–19.

[16] Klaus Krippendorff. 2004. Content analysis: An introduction to its methodology. Sage.

[17] Hao-Ping Lee, Kuan-Yin Chen, Chih-Heng Lin, Chia-Yu Chen, Yu-Lin Chung, Yung-Ju Chang, and Chien-Ru Sun. 2019. Does Who Matter?: Studying the Impact of Relationship Characteristics on Receptivity to Mobile IM Messages. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19), 526:1–526:12. DOI: https://doi.org/10.1145/3290605.3300756

[18] Richard Ling and Birgitte Yttri. 2002. 10 Hyper–coordination via mobile phones in Norway. Perpetual contact: Mobile communication, private talk, public performance 139, (2002).

[19] Afra Mashhadi, Akhil Mathur, and Fahim Kawsar. 2014. The Myth of Subtle Notifications. In Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication (UbiComp '14 Adjunct), 111–114. DOI:h ttps://doi.org/10.1145/2638728.2638759

[20] Abhinav Mehrotra, Robert Hendley, and Mirco Musolesi. 2016. PrefMiner: Mining User's Preferences for Intelligent Mobile Notification Management. In Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16), 1223–1234. DOI: https://doi.org/10.1145/2971648.2971747

[21] Abhinav Mehrotra, Mirco Musolesi, Robert Hendley, and Veljko Pejovic. 2015. Designing Content-driven Intelligent Notification Mechanisms for Mobile Applications. In Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15), 813–824. DOI: https://doi.org/10.1145/2750858.2807544

[22] Abhinav Mehrotra, Veljko Pejovic, Jo Vermeulen, Robert Hendley, and Mirco Musolesi. 2016. My Phone and Me: Understanding People's Receptivity to Mobile Notifications. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16), 1021–1032. DOI: https://doi.org/10.1145/2858036.2858566

[23] Abhinav Mehrotra, Fani Tsapeli, Robert Hendley, and Mirco Musolesi. 2017. MyTraces: Investigating Correlation and Causation Between Users' Emotional States and Mobile Phone Interaction. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. 1, 3 (September 2017), 83:1–83:21. DOI: https://doi.org/10.1145/3130948

[24] Midas Nouwens, Carla F. Griggio, and Wendy E. Mackay. 2017. "WhatsApp is for Family; Messenger is for Friends": Communication Places in App Ecosystems. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17), 727–735. DOI: https://doi.org/10.1145/3025453.3025484

[25] Chunjong Park, Junsung Lim, Juho Kim, Sung-Ju Lee, and Dongman Lee. 2017. Don'T Bother Me. I'M Socializing!: A Breakpoint-Based Smartphone Notification System. In Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17), 541–554. DOI: https://doi.org/10.1145/2998181.2998189

[26] Veljko Pejovic and Mirco Musolesi. 2014. InterruptMe: Designing Intelligent Prompting Mechanisms for Pervasive Applications. In Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '14), 897–908. DOI: https://doi.org/10.1145/2632048.2632062

[27] Martin Pielot, Karen Church, and Rodrigo de Oliveira. 2014. An In-situ Study of Mobile Phone Notifications. In Proceedings of the 16th International Conference on Human-computer Interaction with Mobile Devices & Services (MobileHCI '14), 233–242. DOI: https://doi.org/10.1145/2628363.2628364

[28] Martin Pielot and Luz Rello. 2015. The Do Not Disturb Challenge: A Day Without Notifications. In Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '15), 1761–1766. DOI: https://doi.org/10.1145/2702613.2732704

[29] Martin Pielot and Luz Rello. 2016. Productive, Anxious, Lonely - 24 Hours Without Push Notifications. arXiv:1612.02314 [cs] (December 2016).

[30] Martin Pielot, Amalia Vradi, and Souneil Park. Dismissed! A Detailed Exploration of How Mobile Phone Users Handle Push Notifications. 11.

[31] B. Poppinga, W. Heuten, and S. Boll. 2014. Sensor-Based Identification of Opportune Moments for Triggering Notifications. IEEE Pervasive Computing 13, 1 (January 2014), 22–29. DOI: https://doi.org/10.1109/MPRV.2014.15

[32] Sami Ronkainen. 2001. Earcons in motion–defining language for an intelligent mobile device. In Proceedings of the International Conference on Auditory Display (ICAD).

[33] Stephanie Rosenthal, Anind K. Dey, and Manuela Veloso. 2011. Using Decision-Theoretic Experience Sampling to Build Personalized Mobile Phone Interruption Models. In Pervasive Computing, Kent Lyons, Jeffrey Hightower and Elaine M. Huang (eds.). Springer Berlin Heidelberg, 170–187.

[34] D. M. Russell and M. Oren. 2009. Retrospective Cued Recall: A Method for Accurately Recalling Previous User Behaviors. In 2009 42nd Hawaii International Conference on System Sciences, 1–9. DOI: https://doi.org/10.1109/HICSS.2009.370

[35] Alireza Sahami Shirazi, Niels Henze, Tilman Dingler, Martin Pielot, Dominik Weber, and Albrecht Schmidt. 2014. Large-scale Assessment of Mobile Notifications. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14), 3055–3064. DOI: https://doi.org/10.1145/2556288.2557189

[36] Hillol Sarker, Moushumi Sharmin, Amin Ahsan Ali, Md. Mahbubur Rahman, Rummana Bari, Syed Monowar Hossain, and Santosh Kumar. 2014. Assessing the Availability of Users to Engage in Just-in-time Intervention in the Natural Environment. In Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '14), 909–920. DOI: https://doi.org/10.1145/2632048.2636082

[37] Florian Schulze and Georg Groh. 2014. Studying How Character of Conversation Affects Personal Receptivity to Mobile Notifications. In Proceedings of the Extended Abstracts of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (CHI EA '14), 1729–1734. DOI: https://doi.org/10.1145/2559206.2581320

[38] Florian Schulze and Georg Groh. 2016. Conversational Context Helps Improve Mobile Notification Management. In Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '16), 518–528. DOI: https://doi.org/10.1145/2935334.2935347

[39] Alireza Sahami Shirazi and Niels Henze. 2015. Assessment of Notifications on Smartwatches. In Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct (MobileHCI '15), 1111–1116. DOI: https://doi.org/10.1145/2786567.2794338

[40] Liam D. Turner, Stuart M. Allen, and Roger M. Whitaker. 2015. Push or Delay? Decomposing Smartphone Notification Response

Behaviour. In Human Behavior Understanding. Springer, Cham, 69–83. DOI: https://doi.org/10.1007/978-3-319-24195-1_6

[41] Liam D. Turner, Stuart M. Allen, and Roger M. Whitaker. 2015. Push or Delay? Decomposing Smartphone Notification Response Behaviour. In Human Behavior Understanding, Albert Ali Salah, Ben J. A. Kröse and Diane J. Cook (eds.). Springer International Publishing, 69–83. DOI: https://doi.org/10.1007/978-3-319-24195-1_6

[42] Aku Visuri, Niels van Berkel, Chu Luo, Jorge Goncalves, Denzil Ferreira, and Vassilis Kostakos. 2017. Challenges of Quantified-self: Encouraging Self-reported Data Logging During Recurrent Smartphone Usage. In Proceedings of the 31st British Computer Society Human Computer Interaction Conference (HCI '17), 81:1–81:7. DOI: https://doi.org/10.14236/ewic/HCI2017.81

[43] Dominik Weber, Alireza Sahami Shirazi, and Niels Henze. 2015. Towards Smart Notifications Using Research in the Large. In Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct (MobileHCI '15), 1117–1122. DOI: https://doi.org/10.1145/2786567.2794334

[44] Dominik Weber, Alexandra Voit, Jonas Auda, Stefan Schneegass, and Niels Henze. 2018. Snooze! Investigating the User-Defined Deferral of Mobile Notifications. (2018).

[45] Tilo Westermann. 2017. Global Analysis of User Behavior in the Field. In User Acceptance of Mobile Notifications. Springer, Singapore, 63–79. DOI: https://doi.org/10.1007/978-981-10-3851-8_4

[46] Fengpeng Yuan, Xianyi Gao, and Janne Lindqvist. 2017. How Busy Are You?: Predicting the Interruptibility Intensity of Mobile Users. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17), 5346–5360. DOI: https://doi.org/10.1145/3025453.3025946