Overcoming Compliance Bias in Self-Report Studies: A Cross-Study Analysis

Niels van Berkel, The University of Melbourne
Jorge Goncalves, The University of Melbourne
Simo Hosio, University of Oulu
Zhanna Sarsenbayeva, The University of Melbourne
Eduardo Velloso, The University of Melbourne
Vassilis Kostakos, The University of Melbourne

ABSTRACT

A popular methodology used for *in situ* observations is the Experience Sampling Method (ESM), in which participants intermittently answer short questionnaires. We analyse a set of recent ESM studies and find substantial differences in the number of collected responses across participants. These differences amount to 'compliance bias', as the experiences of responsive participants skew the results. Our work develops ways for researchers to ensure the collection of an adequate number of responses across participants. Through a cross-study analysis of ESM studies, we construct a model that describes the effect of contextual, routine, and study-specific factors on participants' response rate. In addition to previous work, which aims to maximise the number of total responses, this work also aims to achieve a more equal distribution of responses between participants. In order to achieve this goal, we analyse which contextual cues can be personalised to achieve a higher response rate. Our results highlight a number of factors that have a strong effect on participants' response rate and can guide the design of future experiments.

KEYWORDS

Experience Sampling Method; Ecological Momentary Assessment; ESM; response rate; completion rate; compliance; self-report; bias.

1 Introduction

Smartphones have become increasingly popular research tools in the HCI community, thanks to their widespread popularity and powerful sensing capabilities [37]. Established methods are now being adapted into digital versions to be used with smartphones. One such method is the Experience Sampling Method (ESM), used to collect participant self-report data *in situ* on, for example, experiences or emotions [24]. In the ESM, participants are asked to answer multiple sets of questions throughout the day, resulting in a rich record of the participant's life regarding the phenomena of interest to the researcher. Therefore, response rate—the percentage of answered questions—is an important metric in ESM studies [43]. A high response rate ensures a wide spread of results across space and time, increasing the ecological validity of the study results [18], "*the occurrence and distribution of stimulus variables in the natural or customary habitat of an individual*" [18]. Due to the importance of achieving high response rates, researchers have explored several ways to motivate participants. Recent examples include visualising response rates to participants [20], gamification [44], psychological empowerment [13], and use of micro-incentives per individual questionnaires [32].

In this paper, we explore the effect of contextual factors on participant response rates in ESM studies. First, we present an analysis of previous experiments that together highlight that the response rates of participants in studies that employ the ESM vary substantially. This is similar to what has been dubbed as *compliance bias* in the Health Sciences, and has been observed in studies on adherence to therapeutic protocols [38]. The challenge lies in that the differences in response rates cause more responsive participants to have a larger impact on the conclusions of a study. In essence, we show that researchers should expect their results to be biased by default, since response rates can vary substantially across participants. Second, we explore what causes participants to have reduced response rates and suggest three approaches to mitigate compliance bias in Experience Sampling studies.

To increase our understanding of (non-)responsive participants, we set out to establish the effect of smartphone usage on response rates. Relatively little work has considered how contextual and routine smartphone usage factors affect participant response rates. An important reason for this is that studies are typically carried out with the goal of investigating a well-

Table 1. Overview of prior contributions on bias in ESM

Bias	Effect
Design bias	Potential bias introduced by the study design. Lahtia *et al.* [25] identify differences in the probability of sampling across different contexts (time- and sensor-based) and find differences in frequency and skewness of data collection triggers.
Novelty bias	Change in behaviour of participant following the introduction of an unknown device [43]. Given today's prevalence of smartphones, usage of participant devices has become more popular.
Observation effect (Hawthorne effect)	Participants may alter their behaviour as they realise they are being observed. Raento *et al.* [37] mention how a participant may alter their phone conversations after being reminded that their conversation is recorded at the onset of their call.
Self-selection bias	Participants who sign up for ESM studies may differ from those who do not, although the effect appears limited [15, 31].
Selective nonresponse bias	Participants may choose not to respond if they feel uncomfortable sharing or reflecting on their current experience [6].

defined phenomenon, and that the study design naturally stems from this research question. Participants are assumed to comply, when in reality the design of the instrument (*e.g.*, questionnaire items) itself affects compliance, as indicated by previous work in survey research [28]. Further, the literature describes several study-specific constructs which may influence response rate, including the level of rapport between researcher and participant, and the participants' intrinsic motivation [24]. This makes it challenging to determine whether a participant's decision to respond to an ESM questionnaire is due to the participant's context or due to study-specific constraints. Though recognising the context based on sensor data might help to distinguish these differences, inferring participant context is challenging due to development costs and lack of specialised skills and is therefore often avoided [37]. Only through a synthesis of multiple studies it is possible to untangle the effect of contextual and routine factors on participant response rates. By adopting a modelling approach, we identify some of the key factors that affect participants' response rates in ESM studies and discuss how the factors manifest.

For these reasons, we have re-analysed data from recent and independent studies within the domain of HCI. Our cross-study analysis models the differences and similarities between these studies, and identifies which factors affect response rates. We found a number of contextual factors (*e.g.*, phone use, screen state) which have a large effect on response rate. Further, we found considerable differences in the effect of contextual factors across participants. Based on our findings, HCI researchers can optimize the scheduling of data collection based on individual participant context. This has the potential to collect data more evenly across participants and ultimately decrease compliance bias.

2 Related Work

Since its introduction as a research method in the 1970's, the Experience Sampling Method (ESM) has continuously evolved alongside technological developments. Over time, participant-owned smartphones have become the *de facto* research instrument for ESM studies. Several researchers state that the use of the participant's personal device reduces the novelty effect [37], and allows for a more natural *in situ* observation [43]. Despite the increased complexity of the research instrument, not much literature exists on the effect of the device's context on study results. In this work we focus specifically on the effect of smartphone context on (differences in) response rate.

2.1 Bias in Experience Sampling

The ESM and similar self-report methods (*e.g.*, Ecological Momentary Assessment (EMA)) were introduced to achieve two main objectives [24]. First, a shift from collecting participant data in the laboratory to data collection in the real world. The underlying motivation is to collect data in a more realistic setting, preserving the ecological validity of the study: "*imagery evoked in laboratory studies is not necessarily typical of experience encountered in real-life situations*" [24]. Second, researchers sought to reduce the reliance on the participants' ability to recall their experiences [43]. Systematic errors in reasoning, recall, and judgement—known as cognitive biases—negatively impact the validity of collected self-report data [21].

These conceptual improvements over previous methods aimed to increase data reliability. Since then, several studies and commentaries have identified biases in the use of the ESM, potentially affecting the reliability of study results if not adequately identified and handled during study design and data analysis. As with any study, biases can be introduced at any stage of the study procedure. Several works have attempted to identify the scale of biases specific to the use of the ESM and

have sometimes proposed solutions to circumvent the introduced problem. Table 1 provides a summarized overview of these biases.

Despite the amount of work on improving study-wide response rates, no work has developed approaches to homogenise response rates across individual participants. As such, compliance bias is understudied in ESM—this work sets out to address this bias.

2.2 Improving Participant Response Rates

In the methodologically related diary method, monetary incentives have been shown to increase response rate [28]. Musthag *et al.* [32] explored the use of micro-incentives in ESM studies. They found that the amount of the offered micro-incentive did not affect response rate. However, their study did not compare response rates using micro-incentives against a typical 'bulk payment' provided at the end of the experiment.

Several researchers have made suggestions on how to interact with participants to ensure sufficient response rates, usually derived from first-hand experience rather than controlled experimentation. For example, Larson & Csikszentmihalyi [24] encourage constructing a 'viable research alliance', where the participant is aware of the importance of the study and their role in data collection. Stone *et al.* [40] suggest researchers to contact participants who have missed multiple days of data contributions to encourage them to (re-)commence data contribution.

Given the importance of the quantity of answers in ESM studies, HCI research has explored new techniques to increase study response rates. Hsieh *et al.* [20] demonstrated how displaying participant responses rates directly to participants led to increased response rates in a desktop-based ESM study. Church *et al.* [8] allowed participants to specify the time of day at which they were able to respond to questionnaires—reducing the intrusiveness of the method. This does however possibly introduce (unwanted) bias as we do not capture data at times which are inconvenient to participants but of interest to the study. Van Berkel *et al.* [44] introduced gamification elements in a smartphone-based study and found that participants in the gamified condition contributed significantly more self-reports. However, a potential downside of this approach is the fact that gamification is not suitable for all study subjects or participant demographics [34]. Zhang *et al.* [52] explored answering questionnaires through 'unlock journaling' (answer a question while unlocking the phone). This method led to a higher frequency of reporting as well as a decreased intrusiveness. A limitation of this method is the single question offered per questionnaire, as well as the input constraints connected to gesture-based unlock input (*e.g.*, no typing). Hernandez *et al.* [17] compared the presentation of questionnaires on smartwatches and head-worn devices (Google Glass) to smartphones. They found a 13% increase in response rate on both smartwatches and head-worn devices as compared to smartphones, although the difference was not statistically significant. Further, the study contained a different number of prompts between conditions [17], making direct comparison challenging. Intille *et al.* [22] presented questionnaires through smartwatch-based '*microinteractions*', resulting in increased response rates – again, for both studies, the limited screen real estate hampers possibilities for questionnaire items. Finally, both Pejovic & Musolesi, and Mehrotra et al. [42, 48] explored the use of an intelligent notification system which interrupts the user based on user context. The system resulted in reduced response times, and more favourably received notifications. Our work bears resemblance to previous work [30, 35]. We are interested in identifying contextual variables that have an effect of response rate. Unlike previous work, however, we aim to identify these effects across *multiple* studies, not just any single study. As we describe next, previous work has linked a number of contextual variables to response rate in ESM studies.

2.4 Effect of Context on Response Rate

Smartphones are devices that are mobile by design, available to be used anywhere and anytime. As such, the context in which these devices are used constantly changes. In fact, the context of use is so critical that it affects the way in which smartphones are used [41]. Pielot *et al.* [36] showed how several features related to recent phone usage (*e.g.*, screen state, last application launch) can work as predictors for notification interaction. Contextual factors such as location, mood, and time have been shown to influence mobile browsing needs and behaviour [26]. These contextual factors exist 'outside' of a user's smartphone, as they are either based on the user's personal or environmental context. Of particular interest for this study are contextual factors that manifest themselves *through* the user's smartphone. For example, incoming notifications, or a battery warning are all outside of the direct control of the user, yet are likely to influence smartphone usage. In the same way that phone usage is affected by these contextual factors, a participant's willingness to respond to ESM questionnaires may be affected by their smartphone context.

We analysed the literature for contextual and behavioural predictors shown to affect smartphone usage and identified whether they have been investigated in the context of the ESM. We summarise these results in Table 2. Our overview indicates that the effect of the majority of aspects shown to influence *phone usage* has not been tested in the context of the ESM. Furthermore, as these variables have been independently identified in mostly separate studies, the relative importance and impact of each one remains unclear. That is why in our work we analyse data from multiple independent studies and try to identify the relative importance of each contextual variable.

Table 2. Literature overview of effect of several smartphone usage aspects on smartphone usage and ESM responses

Aspects	Phone usage	Effect on ESM behaviour
Battery level	"*[users] limit device use for 'emergency situations'*" [42]	No information
Recent notification history	"[notifications] lead to a reduction in work productivity, including the resumption time from the interruption back to the primary task and the quality and amount of time available for decision making" [34]	No information
Recent app. usage	"a mobile user is likely to use different applications at different locations and access different websites at different times of the day" [39]	No information
Last phone usage	"the majority of phone usage sessions that contain a gap (i.e., phone went to standby mode) consist of only one additional continuous session" [46]	No information
Time of day (diurnal pattern)	"*Roughly 70% of the users in each dataset have a peak hour usage that is more than twice their mean usage.*" [11] "*people are attentive to messages 12.1 hours of the day, [...], people are more attentive during the evening*" [9]	"*we observe that most participants respond across working hours, and late night / early morning hours are less active*" [46]
Weekday / Weekend	"*attentiveness is higher during the week than on the weekend*" [9]	"*Sunday responses were fewer and had a greater variance in happiness*" [5]
Daily apps.	"*users exhibit extraordinarily diverse usage patterns*" [39]	No information
Daily notif.	"[...] sentiment towards a notification varies with the type, completion level and complexity of an ongoing task and the recipient's relationship with the sender." [29]	No information

3 Data Sets

To ensure consistency within our results, we selected three recent studies published in the HCI community which collected contextual data of both answered and unanswered questionnaires. All three studies use the AWARE framework [12] for collecting smartphone sensor data, ensuring consistency in data collection and formatting. Following data cleaning, our combined dataset consists of 8370 ESM notifications (4408, 1649, and 2313 notifications for Dataset 1, 2, and 3 respectively). In all studies, participants used their personal Android smartphones, providing realistic contextual data on their smartphone usage. Across all studies, participants were full-time students from a single University. Although it can be argued that this biases our results, it simultaneously ensures that we control for *e.g.* cultural differences and shifts in power dimensions between researchers and participants which would occur with a more varied sample of studies. We therefore consider this a reasonable trade-off between participant diversity and the control of external variables. We provide a summary of the three study designs and their respective ESM questionnaires below.

Dataset 1: Usage sessions [46]. This study analysed the intentions of smartphone users upon unlocking their device. The goal of the study was to quantify the commonly held assumption that users who lock and unlock their phone shortly after are continuing the same "session" of interaction. To achieve this goal, the authors employed a combination of human sensing (detailed below) and active sensing (*i.e.*, the collection of smartphone sensor data – most notably phone usage). The sample consisted of 17 participants (13 male, 4 female), average age 26. In this seven-day field study, participants were presented with a pop-up ESM questionnaire each time they unlocked their device. The questionnaire consisted of a single question regarding their phone usage ("*Why did you start using your phone?*") with two predefined answer buttons ("*Continue previous objective*" and "*Start on a new objective*"). Upon completion of the study, participants received a fixed reward (movie voucher) as compensation.

Dataset 2: Battery value [19]. This study examined the monetary value that smartphone users assign to their device's battery life, dependent on their context, needs, and current battery level. To quantify the monetary value held by participants, the study incorporated a reverse second-price sealed-bid auction protocol. Upon receiving an ESM question ("*How much money should we pay you for 10% units of your battery?*"), participants bid their desired amount of money for the exchange of 10% of their smartphone's current battery life. Using the aforementioned auction protocol, the participant with the lowest bid won the auction and received the money offered in the second-lowest bid after draining their battery by 10%. The study lasted for eight days and consisted of a total of 20 participants (17 male, 5 female) with an average age of 24. Participants received 13 notifications per day following an hourly schedule (interval contingent). ESM notifications expired after 10 minutes. Participants received a fixed reward (€50) as well as any money won in the auction.

Dataset 3: Scheduling effects [45]. This study assessed the effect of different ESM schedules on participant response rate and recall accuracy. Participants answered questionnaires consisting of five questions, all focused on assessing their smartphone usage (*e.g.*, "*How many unique applications have you used since 12:00?*"). We consider only the initial question of each questionnaire as opposed to all individual questions (which would register as identical). Participants received up to six notifications a day, with a weekly-changing within-subjects design determining the timing of notifications; randomised (signal contingent), every other hour (interval contingent), or upon unlocking their phone (event contingent, but restricted to one per time-period). ESM notifications were set to expire after 15 minutes. Data was collected over a period of 3 weeks, with a total of 20 participants (15 male, 5 female), average age of 26. Following completion of the study, participants received a fixed reward (two movie vouchers).

## 4   Analysis I – Compliance Bias in ESM

In Analysis I we investigate the existence and magnitude of compliance bias in ESM. We define compliance bias as the introduction of systematic predilection into collected self-reports as the result of differences in response rate between participants. Because participant responses are typically analysed as a single entity, over- and under-representation of participants in the collected dataset can significantly bias the analysis of results. The literature shows that differences in the number of participant responses are not uncommon. Wiese et al. [48] found that their most responsive participant completed 358 questionnaires, whereas their least responsive participant completed 29 questionnaires. Similar results are presented by Epp et al. [10], reporting participant responses ranging from 2 to 219.

For our analysis of compliance bias, in addition to the three previously introduced studies, we also analyse the data from the *StudentLife* study [47], a publicly available dataset containing ESM responses. The StudentLife dataset was collected over a 10-week term from 48 University students, assessing their mental health, academic performance, and behaviour. The dataset contains various questionnaires, each triggered according to a different schedule. For consistency, we restrict our analysis to only one of these questionnaires. We select the questionnaire for which most responses were collected. Participants were asked to indicate their current state by selecting a photo amongst a grid of photos depicting various emotions (Photographic Affect Meter). Since this dataset contains only answered ESM questionnaires (no unanswered questionnaires), we are unable to calculate individual response rates per participant and instead consider the absolute number of participant responses.

4.1   Results and Discussion

We calculate and visualise the number of responses as collected from individual participants for all four studies (Figure 1). Substantial differences between individual participants exist in all four studies. This is also indicated by the large standard deviation in the collected number of responses between participants. All four studies are remarkably consistent in the fact that 30% of participants provide roughly 50% of the total collected responses. The 30% of participants with the lowest number of responses account for merely 6—17% of responses.

These results indicate that there is serious compliance bias in the four analysed studies, despite the fact that aggregated data is used to infer study results [46]. Although these numbers are rarely reported in the literature, given these results and the stark differences reported by Wiese *et al.* [48] and Epp *et al.* [10], we have reasons to believe that compliance bias may be widespread in self-report studies. Therefore, the current approach of data analysis, in which all responses are considered

Dataset 1. Usage sessions [46].
Resp./participant: 243.1 (±127.4)

Dataset 2. Battery value [19].
Resp./participant: 44.1 (±29.2)

Dataset 3. Scheduling effects [45].
Resp./participant: 69.0 (±26.3)

StudentLife [47].
Resp./participant: 184.5 (±107.1)

Figure 1. Number of responses per individual participant across four different studies.

equal, causes responsive participants to have a larger effect on study conclusions. Following the identified differences between participants, we now present a more detailed analysis of variables leading to participant non-response.

5 Analysis II – Modelling Compliance Bias

Whereas we establish the existence of compliance bias in Analysis I, we now construct a more detailed insight of the effect of (smartphone-based) contextual, routine, and study-specific characteristics on response rate. Establishing the effect of these characteristics could assist researchers in designing more intelligent notification schemes or encourage future studies into the effect of these contextual factors on study participants.

5.1 Method

We combine the three aforementioned datasets into one dataset to investigate compliance bias more systematically than possible in a single study. Here, each row consists of a single ESM question (binary: answered or unanswered). We exclude the StudentLife dataset (included in Analysis I) in this analysis as the dataset is restricted to answered questionnaires only, missing the critical context of unanswered questionnaires.

In addition to participant responses, we collected contextual data in each of the three studies (detailed below). We removed elements in our dataset that were incorrectly marked as a notification (*e.g.*, applications which use notifications to display audio controls, resulting in a large number of incorrect 'notifications'). Data was recorded upon a state change for each respective data element, combined with both a unique random ID per participant and a timestamp. From the contextual usage data contained in our dataset, we calculated the following predictors:

Contextual predictors. Contextual predictors reflect the *in situ* participant state at the moment of ESM notification arrival. These predictors are constantly changing, and jointly affected by smartphone usage and the timing of ESM notifications.

- **Battery level**: Percentage of remaining battery life.
- **Charging status**: Smartphone's charging status, either 'charging' or 'not charging'.
- **Screen status**: State of the screen, either 'on' or 'off'.
- **Last phone usage**: Time in minutes since the phone was last used. Values above 60 minutes are clamped to 60 minutes.
- **Recent notifications**. Number of notifications received by the participant in the preceding 15 minutes.
- **Weekend/weekday**. We assigned either 'weekend' or 'weekday' based on the day the ESM notification was sent.
- **Time of day**. Questionnaire arrival in morning [05:00 - 11:59], afternoon [12:00 - 16:59], or evening [17:00 - 04:59].

Routine predictors. Routine predictors characterise the smartphone usage behaviour across the duration of the study.

- **Daily application usage**. Number of average unique applications a participant uses per day.
- **Daily notifications**. Number of average notifications a participant receives per day.

Study-specific characteristics. We specify one key ESM study configuration parameter, daily ESM notifications, which is not identical between the three studies. Other study-specific characteristics (*e.g.*, questionnaire input type, study duration) are not suitable to be included in the model. Though we expect these parameters to affect a participant's willingness to respond, including these parameters in the model would reduce its reliability. If included, these parameters would assess all potential differences between the studies rather than one single variable. For example, using 'input type' as a parameter, its value would be 'popup' for Dataset 1 and 'notification' for the two remaining datasets. As such, the effects ascribed to this coefficient would model any differences between the studies rather than actually measuring the effect of input type.

- **Daily ESM notifications**. Average daily ESM notifications per dataset are 41.2 (± 18.39), 11.0 (± 1.49), and 5.4 (± 0.28) for Dataset 1, 2, and 3 respectively.

Criteria were selected based on recommendations from the related work as well as their shared availability across datasets. Rows for which one of the above predictors were missing were removed from the dataset to ensure reliable model construction. Following this, 8185 rows of ESM questions remained (185 were removed).

6 Results

Given our binary predictor value of ESM responses—answered or unanswered—the distribution of participant responses is binomial. Rather than considering all of the collected data to originate from one coherent source, we have to account for the fact that both the number of observations per study as well as the number of observations per participant differ from one another. We therefore construct a generalised linear mixed-effects model (GLMM) with a binomial distribution describing the effect of the aforementioned predictors on compliance. We specify both 'participant ID' and 'study ID' as random effects in the model. As a result, our model treats participant- and study-specific variations as a nuisance term (*i.e.*, it is not of direct interest but should be controlled for). The use of generalised linear-mixed effects models for binary responses is an increasingly common paradigm applied to data collected in a longitudinal setting [51]. We create one general model based on the data of all participants. Then, we create separate models for the 30% of participants with the highest response rate (per individual study, roughly half of the collected responses) and the 30% participants with the lowest response rate (again per individual study). We label these two separate groups as high-compliance and low-compliance.

Table 3. Confusion matrices for the three models.

	General model		High-compliance		Low-compliance	
	Predicted		Predicted		Predicted	
	Unansw.	Answ.	Unansw.	Answ.	Unansw.	Answ.
Unasw.	998	1190	14	230	705	398
Answ.	440	5557	8	2232	255	1070
Accuracy	80.1%		90.4%		73.1%	
Sensitivity	92.7%		99.6%		80.8%	
Specificity	45.6%		5.7%		63.9%	
Baseline	73.3%		90.2%		54.6%	

Predictors that were not significant ($p < .05$) across any of the three models were discarded. This was true for two variables: charging status and weekend/weekday. The three models were recalculated excluding these two variables. We report the outcomes of each model (including odds ratio and confidence interval per predictor) in Table 4. The general model is statistically significant (χ^2 (12) = 354.5, $p < .001$) and explains 8% of variance in answering ESM notifications (*Marginal R^2* = .08, *Conditional R^2* = .38). The model for the most responsive participants is also statistically significant (χ^2 (12) = 97.1, $p < .001$) and explained 15% of variance in answering ESM notifications (*Marginal R^2* = 0.15, *Conditional R^2* = 0.28). Finally, the model for the least responsive participants is also statistically significant (χ^2 (12) = 141.8, $p < .001$) and explained 28% of variance in answering ESM notifications (*Marginal R^2* = 0.28, *Conditional R^2* = 0.39).

When assessing human behaviour we do not expect our model to include all relevant predictors [3]. R^2 describes the goodness-of-fit of the model and was calculated using the R package '*piecewiseSEM*' [27], as based on the method introduced in [33] specifically aimed at GLMM. Marginal R^2 describes the variance explained solely by fixed factors, whereas the Conditional R^2 describes the variance explained by the combined fixed and random factors (in our case 'participant ID' and 'study ID') [33].

We present the confusion matrices for our classification models in Table 3. Further, Table 3 reports the accuracy of the three models as well as their respective sensitivity and specificity values. Sensitivity describes the percentage of cases in which the model predicts 'Answered' among actual 'Answered' cases (also described as 'recall' or 'true positive rate'). Specificity describes the percentage of cases in which the model predicts 'Unanswered' among the total number of actual 'Unanswered' cases.

Table 4. Effect of model factors on response rate (odds ratio, confidence interval, and p-value).

	General model			High-compliance			Low-compliance		
	OR	CI	p	OR	CI	p	OR	CI	p
(Intercept)	7.52	2.0-28.0	<.003**	15.26	3.2-72.1	<.001***	0.6	0.2-1.6	.274
Screen status – On	0.40	0.3-0.5	<.001***	0.20	0.1-0.3	<.001***	0.4	0.3-0.5	<.001***
Time of day – Evening	0.79	0.7-0.9	.004**	1.54	1.0-2.3	.033*	0.8	0.6-1.0	.079
Time of day – Afternoon	0.93	0.8-1.1	.362	1.33	0.9-1.9	.128	0.9	0.7-1.2	.510
Last phone usage	0.98	1.0-1.0	<.001***	0.98	1.0-1.0	<.001***	1.0	1.0-1.0	<.001***
Recent notifications	0.99	1.0-1.0	.048*	1.03	1.0-1.1	.435	1.0	1.0-1.0	.501
Battery level	1.00	1.0-1.0	.012*	1.00	1.0-1.0	.353	1.0	1.0-1.0	.836
Daily notifications	1.00	1.0-1.0	.799	1.00	1.0-1.0	.970	1.0	1.0-1.0	.049*
Daily ESM notifications	1.01	0.9-1.1	.392	1.04	1.0-1.1	.092	1.0	1.0-1.1	.042*
Daily application usage	1.06	1.0-1.2	.240	1.08	0.9-1.2	.271	1.2	1.1-1.4	.003**

To ensure that our chosen predictors are not in fact measuring the same phenomenon, we check for the existence of multicollinearity. Multicollinearity indicates whether one of the predictors in the model can be linearly predicted from a combination of any of the other predictors, reducing the validity of individual predictors. All our predictors report a variance inflation factor between 1.01 and 1.73, well below the often-used threshold of 5 or 10 to detect (severe) multicollinearity [14]. We also check for linearity between predictors and dependent variable in the three constructed models. The analysis conducted with the '*caret*' package [23] reveals no linear combinations between any of the predictors and the dependent variable.

We plot the effect of each predictor in the model in Figure 2. This shows the effect of, for example, the smartphone's screen status on ESM response rate, assuming that all other predictors remain unchanged. We note that the Y-axis is not consistent between graphs: its range is, in essence, an indicator of the effect size for each predictor. As shown in Figure 2, we observe that:

- Participants are more likely to respond when they receive a notification while their screen is off; *ergo*, they are not using their phone (stronger for low-compliance participants).
- Participants are more likely to respond in the morning, with a small decline in the transition to afternoon and evening. The opposite holds for high-compliance participants.
- Participants are more likely to respond when they have recently used the phone. This effect applies to all participants but is reduced for high-compliance participants.
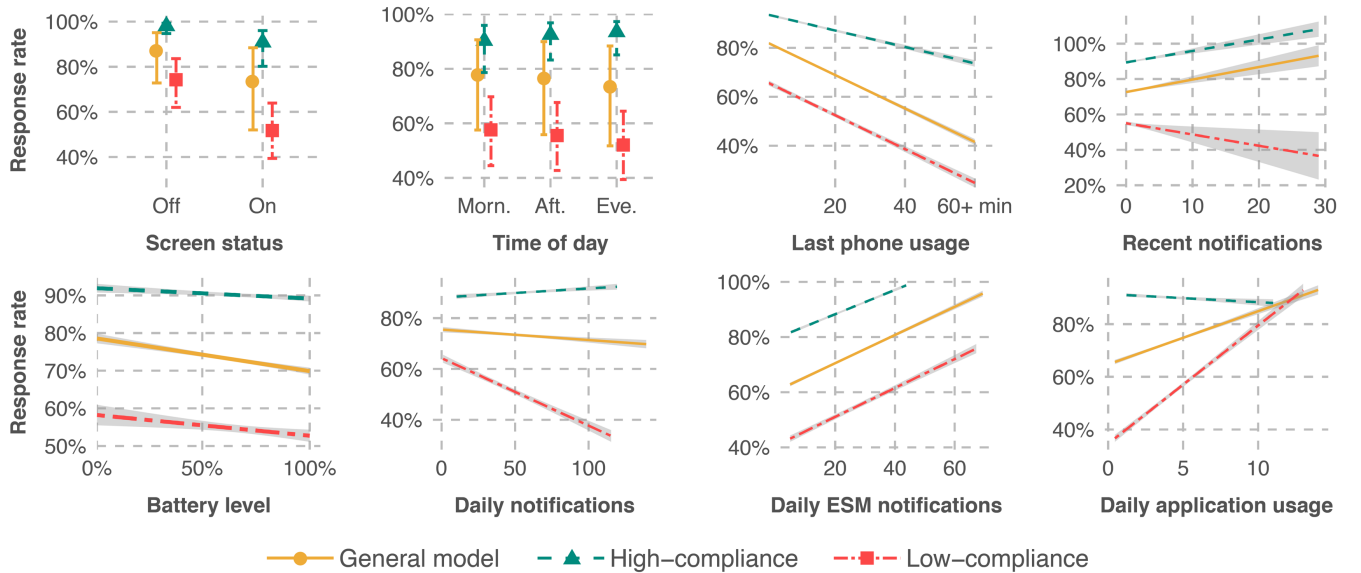
Figure 2. Effect of individual predictors, with all other predictors remaining equal.

- High-compliance participants increase their likelihood of response when experiencing a high number of incoming notifications, whereas low-compliance participants are less likely to respond when receiving many notifications.
- A lower battery level results in a higher response rate. This effect is almost negligible for high-compliance participants.
- Number of daily notifications (averaged over study period) received by participants shows a discrepancy between participant clusters. Low-compliance participants with a high number of daily notifications are less responsive than their peers, whereas the opposite holds for high-compliance participants with a high number of daily notifications.
- More daily notifications lead to a higher response rate.
- In general, participants who use a high number of unique daily applications (averaged over the study period) are more likely to respond to ESM notifications.

Next, we take a more detailed look at a few of these predictors. This is because, in the case of continuous variables, the effect of a predictor is not always as straightforward as portrayed in Figure 2. Therefore, we calculate the density plots for predictors which show an effect of context on response rate.

Effect of Last Phone Usage

We calculated the effect of last phone usage against ESM response rate, as displayed in Figure 3 (excluding the values above 60 min.). A participant who has recently used their device is considerably more likely to answer an ESM notification, whereas a longer time since the last device usage results in a higher percentage of unanswered questionnaires. This effect is strongest for participants with a low response rate.


Figure 3. Conditional density plot of last phone usage and ESM response rate.

Effect of Recent Notifications

As shown in Figure 4, the effect of recent smartphone notifications is non-linear. Participants with either zero or a large number of recent notifications are more likely to answer (given the density of answered over unanswered at these values). This effect is most limited for participants with a low response rate.


Figure 4. Conditional density plot of recent notifications and ESM response rate.

Effect of Battery Level

Figure 5 shows the effect of different battery levels on ESM response rate. For many ranges of battery levels, the chance of an ESM being answered or unanswered is equal. Only as the battery level approaches 100% does the chance of an ESM going unanswered increase. This effect is strongest for high-compliance participants.

Figure 5. Conditional density plot of battery level and ESM response rate.

Effect of Daily Notifications

The effect of the daily number of notifications on response rate shows an inconsistent pattern (Figure 6). We see that for both the general model and the model describing the low-compliance participants, receiving many daily notifications is likely to lead to a reduction in response rate. The opposite holds for high-compliance participants, where a low number of daily notifications actually results in a lower response rate than their (high-responsive) peers.


Figure 6. Conditional density plot of daily notifications and ESM response rate.

Effect of Daily Application Usage

Participants who use a low number of daily applications are less likely to respond than those who use a higher number of applications. As shown in Figure 7, this effect holds across all three clusters, but is most prominent in the general model and the model assessing the most responsive participants. Participants with a low number of daily applications used are maybe more likely to use a secondary device (*e.g.*, laptop, tablet) for their browsing and communication needs, or are simply less likely to use their smartphone throughout their day – thereby failing to notice incoming ESM notifications.


Figure 7. Conditional density plot of daily unique application usage and ESM response rate.

7 Discussion

Researchers employing the ESM rely on their participants' responses to make inferences about the topic being studied. Although a strength of the ESM is the repetitive collection of self-reports over time in the participants' natural environments, the same participant responses also become the *Achilles Heel* of the method. Several biases related to ESM studies have been previously identified, including design bias [25], novelty bias [43], and self-selection bias [15, 31], among others (see Table 1). Further, a low number of participant responses has been discussed as detrimental to the reliability of ESM studies [24]. Recent discussion in Psychology and other disciplines have also critiqued the current bias towards '*WEIRD*' (well-educated

and originating from industrialised, rich, and democratic countries) and university participant samples [16]. Such samples are also common among Experience Sampling studies, as seen in the studies included in our study, as well as most other recent ESM publications (*e.g.*, [47, 50]). These participants are not a realistic representation of society, limiting study conclusions to the studied population.

Our results show that there is a considerable difference in the number of contributions between participants (Figure 1). This phenomenon, which has not previously been systematically summarized and explored, appears to be widespread (*e.g.* [10, 47, 50]). We term this difference in responses between participants 'compliance bias'. To overcome these differences between participants, previous work recommends the removal of participants with a low response rate [43]. For example, Epp *et al.* [10] removed those with less than 50 responses, and Yang *et al.* [50] removed participants who completed less than half of the assigned tasks. However, we argue this approach is problematic and at the very least not optimal. Not only is such a selection of acceptable response numbers arbitrary, it also removes the experiences of a set of participants from the analysed data – biasing results to a smaller and more engaged participant sample. We now interpret the presented results, offer suggestions on reducing compliance bias, and explain how compliance bias relates to previously identified biases in self-report studies.

7.1 Modelling Results

From the total of 10 predictors analysed in this study, 8 were found to be significant in at least one of the models. From this, we conclude that different contextual factors significantly affect participant response rates. Other than time-related contextual variables, the selected predictors have not been previously explored in the context of ESM response rate (see Table 2). Questionnaire scheduling techniques in ESM studies often assume that a participant's ability to respond to a questionnaire is consistent both across sampling time and within the participant sample. Our results, on the contrary, show that participant response rates are not consistent across time and context, and that significant differences exist between participants. For example, we find that study participants are less likely to answer after leaving their phone unused for a long duration, or when their phone is currently being used. We also find that the response rate of high-compliance participants is even higher during the evening, while the opposite holds for low-compliance participants. Similarly, high-compliance participants are more responsive when having just received a higher number of notifications, whereas low-compliance respondents become even less responsive under the same condition.

Contextual Predictors

The context of a participant who receives an ESM questionnaire affects the participant's ability and willingness to reply. This is why, for example, many ESM studies avoid sending notifications during the night [19]. Pejovic & Musolesi [35] identify a variety of contextual factors related to ESM responsiveness, but do not consider smartphone usage context. Our results show that various smartphone usage factors significantly affect response rate.

Participants' *screen status*, either 'on' or 'off', has a significant effect on response rate. Interestingly, we find a higher response rate when the screen of the device is turned off. This indicates that when participants are currently already using their device for a specific purpose, their focus on that task tends to take precedence over answering an ESM notification. In addition, we found that *last phone usage* is also a significant indicator of ESM response rate. We observe that an ESM questionnaire is more likely to be answered if the participant has recently used their phone. A large time gap may indicate that the participant is busy with other activities or simply not in the vicinity of the device. This is in line with previous work on smartphone usage [8, 46], but has not been considered in light of ESM studies. We infer that a participant is more likely to respond to an ESM notification when the phone is on the periphery of attention but not in active use. These results hold for both high- and low-compliance participants.

Time of day, categorised as morning, afternoon, and evening, also had a significant effect on ESM response rate. For most participants, notifications sent in the morning or afternoon hours are more likely to be answered than those in the evening hours. As a result, if a researcher is to send notifications equally distributed over the course of day, the collected response data will be skewed towards the morning and afternoon. The opposite holds for high-responsive participants, who are actually responding slightly more frequently in the evening hours. In their analysis of smartphone user attentiveness to notifications, Dingler & Pielot found that levels of attentiveness are higher during evenings [9]. Our results show that, in contrast, study participants are less responsive to ESM questionnaires during evening hours. This highlights a difference between notifications originating from user-installed applications and questionnaires.

Previous work has identified that notifications are experienced as interrupting [7], unfavourable [35], and annoying following a high frequency of notifications [2]. As such, it can be argued that interrupting participants with a questionnaire following a high number of notifications may lead to frustrated participants and dismissed ESM notifications. Conversely,

incoming notifications may indicate active usage of the device and therefore indicate a suitable moment for the interruption. Our results show that the effect of *recent notification history* differs between high-compliance participants and low-compliance participants. Among low-compliance participants we see a decrease in their responsiveness when receiving a large number of notifications.

Finally, a participant's *smartphone's battery life* was found to have a significant effect on participant response rate. According to the generated model, a higher battery life results in a lower response rate, although the size of the effect is limited. As shown in Figure 5, the effect of battery life is not linear but instead primarily prominent at the 'edge-value' of battery life. The chance of ESM notifications going unanswered is highest in the case of full or almost-full battery life. A battery level of 100% is an indicator that the phone is likely to be connected to a power source and charging. Despite the significant effect of battery level, the effect size is limited (Table 4) – explaining the exclusion of charging status as a predictor.

Earlier work on the ESM recommends participant sampling across both weekdays and weekends, to capture a diverse set of experiences [15]. Csikszentmihalyi *et al.* [5] report fewer participant responses on Sundays. Our results imply that, for our data, no systematic bias was introduced by sampling across both weekdays and weekends. In the three investigated studies, participants used their personal devices, whereas the studies completed by Csikszentmihalyi *et al.* [5] required participants to carry both a beeper and pen-and-paper – making it more likely for participants to forget their recording device or experience it as burdensome.

Routine Predictors
Smartphone usage differs drastically between people. Thus, smartphone users can be categorised by their usage habits [49]. Here, we investigated two routine smartphone usage parameters: application usage and notifications received.

We find that participants that use a larger number of *unique daily applications* are more responsive to ESM notifications. This effect holds true specifically among low-compliance participants. Further, we find differences between high- and low-compliance participants in the effect of *daily notifications*. Among high-compliance participants, response rate is even higher for those with a high number of daily notifications. For low-compliance participants, a high number of daily notifications is an indicator of even lower response rates. Whereas high-compliance participants seem to be able and willing to manage ESM requests among incoming smartphone notifications, the opposite holds for low-compliance participants.

Study-specific Predictors
The effect of methodological configurations in Experience Sampling on response rate has typically been discussed from the perspective of participant strain. For example, Consolvo & Walker [4] suggest designing questionnaires that minimise participants' burden by avoiding open-ended questions or reducing the number of notifications. Mehrotra *et al.* [30] note that "*users may fail to respond honestly, or may even ignore the questionnaire prompts if they perceive the study as too burdensome*".

The sampled studies contain differences in the amount of *daily ESM notifications* presented to participants. The number of daily questionnaires is often a balance between the researchers' desire for a rich data set versus participant burden [53]. Our results show that an increase in the number of ESM notifications leads to an increase in response rate. This effect is highly consistent between all three models. However, we do note that the sampled studies contained differences in questionnaire structure. In the 'Usage Sessions' study, participants received a *popup message* with a binary question upon device unlock. While this may have been experienced as interrupting, the required effort was lower when compared to the two other studies (unlocking the phone and opening the notification). We control for this difference by specifying study (and participant) as random factors in our model. Although there may be a variety of other study-specific predictors causing differences in response rate (*e.g.*, study topic), our data does not allow for analysing the effect of such factors and we therefore do not pose any claims regarding other study-specific factors. Nevertheless, our analysis shows a multitude of factors that do significantly impact participant response rates – explaining up to 28% of variance among low-compliance participants, the most critical participant group when considering compliance bias. Given the nature of our cross-study analysis we can expect these factors to hold for future ESM studies.

7.2 Mitigating Compliance Bias in ESM
Our analysis of multiple recent ESM studies shows that compliance bias exists, and that it has a substantial effect on the interpretation of study results. From the analysed studies, half of contributions were made by 30% of participations, whereas the 30% with the lowest number of contributions were responsible for only 6-17% of responses (see Figure 1). A good first step towards addressing compliance bias in ESM is therefore to recognize that this bias exists, and to urge researchers to analyse and report the difference in collected participant responses.

Compliance bias can be reduced by either oversampling those participants that are likely to be underrepresented in the collected dataset, or by undersampling highly responsive participants. While discarding responses following data collection is the most straightforward approach to undersampling, it is also detrimental to many of the core aspects of the ESM. Typically, the more responses are collected, the richer the analysis that can be performed by the researcher [6]. Therefore, researchers are likely to want to focus on *increasing the amount* of collected responses of low-respondents rather than decreasing responses in an attempt to homogenise response rates. One approach to increase the number of collected responses among low-respondents is to extend the study duration for low-compliance participants. This will result in a higher number of absolute responses, but is often impractical due to research agreements, will delay the analysis, and conflicts with other practical study arrangements. We therefore introduce and discuss three different strategies to homogenise participant responses.

Contextual optimisation. Low-responsive participants are not only more likely to ignore the questionnaire when their device is in active use, but also when it is requesting their attention (recent notifications (Table 4). This shows that high-responsive participants are more willing to overcome these 'barriers' to answer a question or retrieving their smartphone after not having used it for a long time. To overcome this difference, one strategy to balance response rates is to personalise questionnaire scheduling to those contexts in which a participant is likely to respond. In a way, this follows prior work by Church *et al.* [1] who allowed participants to specify the times of day at which they would receive ESM questionnaires. We argue that future work should explore the personalisation of ESM sampling beyond time-restrictions, including contextual factors. These personalised schedules adapt their schedule to interrupt participants at moments in which a response is likely. This builds on earlier recommendations (*e.g.*, [30, 35]). Personalised scheduling based on participant context will result in a more equal response rate between participants, reducing the effect of compliance bias. The results we have presented here provide a starting point for which contextual variables to consider in this approach.

Increase answer opportunities. Participants with a low response rate respond to ESM questionnaires less frequently. To collect an equal number of responses between participants, participants with a lower response rate in comparison to their peers can be sampled more frequently. Naturally, this will mean that the study orchestrators will have to find a way to adjust the rate on the go, while the study is progressing. With the use of mobile phones and modern ESM tools, this is luckily quite feasible already. While this will result in a further decrease of their individual response rate (as well as global study response rate), it allows the researcher to homogenise the absolute number of collected responses between participants. This approach does not consider participant context and could lead to increased annoyance among participants – potentially having a detrimental effect on the quality of collected responses. One way to reduce this annoyance is to oversample only in those contexts where the participants miss the questionnaire (expired questionnaires) rather than actively ignore it (dismisses questionnaires).

Context-based oversampling and undersampling. A different approach to addressing compliance bias is to focus on the breath of contexts covered in collected responses rather than the total number of collected responses. For example, when a participant is biased towards answering questionnaires when they are messaging friends, future questionnaires should undersample during this context, and extra questionnaires should be scheduled during contexts which lack data. While this is likely to result in a lower number of overall responses, the difference in the number of collected responses between participants is likely to be reduced and the ecological validity of the results is increased.

Naturally, these suggestions are not always applicable. For instance, Church *et al.'s* study design [1] prevented the collection of participant responses at moments that participants deemed inconvenient. Similarly, the sampling of experiences at moments at which a participant is more likely to respond will reduce the variety of contexts captured during the study. The effect of this on study results is highly dependent on the study question and the larger goal of the study. When focusing on, for example, a participant's stress levels, it is critical to sample responses across a variety of (smartphone usage) contexts, including when a participant is experiencing a large number of notifications or an alarmingly low battery level. This is why we urge researchers to strongly consider which parameters to include in personalised ESM models based on established literature.

Finally, when addressing compliance bias in self-report studies, it is important to consider the potential interplay with other biases. Attempts at reducing compliance bias can result in the introduction or reinforcement of other biases. First, sampling participants during contexts in which they are likely to be more responsive will limit the number of contexts in which participants answer questionnaires [25]. Simultaneously, smartphone usage behaviour may differ between participants, which could result in participants receiving a significantly different number of ESM questionnaires if following a single notification scheme for all.

Second, the recruitment of university students will likely lead to different smartphone usage patterns than displayed by other samples of the population. Therefore, the use of sampling techniques tested among university student participants may not necessarily be successful among other parts of the population.

Third, novelty bias can affect compliance bias over time. It is well known that participant engagement drops over time, resulting in a reducing number of responses [40]. As a result, study results are likely to be biased towards the initial sampling period where participants still experience the study as interesting.

Our results show that, in addition to study-specific predictors, a combination of contextual and routine smartphone usage predictors as measured from participant smartphones significantly affect ESM response rate. Further, we show that these predictors affect high- and low-compliance participants differently. These findings can be useful for future ESM studies and their question scheduling, as a high response rate is crucial to securing ecological validity of the study [18]. We position our work as one of the contributions of HCI to the continuous development of the ESM—with contributions ranging from novel input methods to more intelligent interruption mechanisms.

7.4 Limitations and Future Work

Our analysis of ESM study data contains several limitations. Given the nature of a cross-study analysis in Analysis II, we are limited to the data collected from the three selected datasets. As a result, we could only include a limited number of methodological study parameters. Even though the publicly available ESM dataset StudentLife [47] contains a wide range of contextual data, the dataset is limited to answered questionnaires and contextual features do not fully overlap. In the presented analysis, the context of both answered and unanswered ESM questionnaires is essential. Future work should identify in more detail the effect of individual contextual and routine parameters on response rate, and how this newfound knowledge can drive notification schedules while retaining ecological validity. The studies included in this analysis range in duration from 7 to 21 days. While this is in line with common ESM practice [43], it is not clear how results would be affected in studies of longer duration.

The population in the three datasets (Analysis II) is limited to the student population of one University. While the similar background of participants ensures that the identified differences are not due to differences in student-researcher power-relationships, motivation, or other cultural aspects, the study sample is limited in variety. The studied population sample (young adults, well-educated) differs from the general population, including in their use of smartphones. We therefore note that our results do not necessarily replicate among other target groups. However, our results are directly applicable to the many researchers relying on student populations and our method can be replicated for future studies with more diverse samples.

8 Conclusion

Through an analysis of four recent ESM studies, including the publicly available StudentLife dataset [47], we identify substantial differences in the number of collected responses between study participants. This introduces a considerable bias in the analysis of study results, as the experiences of high-compliance participants have a disproportionate effect during study analysis. We term this compliance bias. Following the establishment of this phenomenon in ESM studies, we perform a cross-study analysis of three recent studies investigating the effect of contextual, routine, and study-specific factors on participants' response rate. We demonstrate that a number of contextual factors are highly predictive of participant response rate. Prominently, sampling when the phone has recently been used and when the screen is off (indicating that it is no longer being used) is likely to lead to higher response rates. Further, our results show that there are considerable differences in the effect of these factors on responsive and irresponsive participants. We propose different scheduling techniques that could help mitigate compliance bias by homogenising the number of participant responses. Finally, we discuss potential concerns with regards to other biases.

REFERENCES

[1] Karen Church, Mauro Cherubini, and Nuria Oliver. 2014. A Large-scale Study of Daily Information Needs Captured In Situ. *ACM Transactions on Computer-Human Interaction*, 21 (2). 1-46. https://doi.org/10.1145/2552193

[2] Karen Church, and Rodrigo de Oliveira. 2013. What's Up with WhatsApp?: Comparing Mobile Instant Messaging Behaviors with Traditional SMS. *International Conference on Human-computer Interaction with Mobile Devices and Services*, ACM, 352-361. https://doi.org/10.1145/2493190.2493225

[3] Jacob Cohen, Patricia Cohen, Stephen G. West, and Leona S. Aiken. 2002. *Applied Multiple Regression/Correlation Analysis for the Behavioral Sciences.* Routledge.

[4] Sunny Consolvo, and Miriam Walker. 2003. Using the Experience Sampling Method to Evaluate Ubicomp Applications. *IEEE Pervasive Computing*, 2 (2). 24-31. https://doi.org/10.1109/MPRV.2003.1203750

[5] Mihaly Csikszentmihalyi, and Jeremy Hunter. 2003. Happiness in Everyday Life: The Uses of Experience Sampling. *Journal of Happiness Studies*, 4 (2). 185-199. https://doi.org/10.1023/a:1024409732742

[6] Mihaly Csikszentmihalyi, and Reed Larson. 2014. Validity and Reliability of the Experience-Sampling Method. In Csikszentmihalyi, M. ed. *Flow and the Foundations of Positive Psychology*, Springer, 35-54.

[7]   Edward Cutrell, Mary Czerwinski, and Eric Horvitz. 2001. Notification, Disruption, and Memory: Effects of Messaging Interruptions on Memory and Performance. *Human-Computer Interaction - INTERACT*, 263-269.

[8]   Anind K. Dey, K. Wac, Denzil Ferreira, K. Tassini, J. H. Hong, and J. Ramos. 2011. Getting Closer: An Empirical Investigation of the Proximity of User to Their Smart Phones. *International Conference on Ubiquitous Computing*, ACM, 163-172. https://doi.org/10.1145/2030112.2030135

[9]   Tilman Dingler, and Martin Pielot. 2015. I'll Be There for You: Quantifying Attentiveness Towards Mobile Messaging. *17th International Conference on Human-Computer Interaction with Mobile Devices and Services*, ACM, 1-5. https://doi.org/10.1145/2785830.2785840

[10]   Clayton Epp, Michael Lippold, and Regan L. Mandryk. 2011. Identifying Emotional States Using Keystroke Dynamics. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 715-724. https://doi.org/10.1145/1978942.1979046

[11]   Hossein Falaki, Ratul Mahajan, Srikanth Kandula, Dimitrios Lymberopoulos, Ramesh Govindan, and Deborah Estrin. 2010. Diversity in Smartphone Usage. *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services*, ACM, 179-194. https://doi.org/10.1145/1814433.1814453

[12]   Denzil Ferreira, Vassilis Kostakos, and Anind K. Dey. 2015. AWARE: Mobile Context Instrumentation Framework. *Frontiers in ICT*, 2 (6). 1-9. https://doi.org/10.3389/fict.2015.00006

[13]   J. Goncalves, V. Kostakos, E. Karapanos, M. Barreto, T. Camacho, A. Tomasic, and J. Zimmerman. 2014. Citizen Motivation on the Go: The Role of Psychological Empowerment. *Interacting with Computers*, 26 (3). 196-207. https://doi.org/10.1093/iwc/iwt035

[14]   Joseph F. Hair, William C. Black, Barry J. Babin, and Rolph E. Anderson. 2010. *Multivariate Data Analysis*. Pearson.

[15]   Joel M. Hektner, Jennifer A. Schmidt, and Mihaly Csikszentmihalyi. 2007. *Experience Sampling Method: Measuring the Quality of Everyday Life*. Sage.

[16]   Joseph Henrich, Steven J Heine, and Ara Norenzayan. 2010. Most people are not WEIRD. *Nature*, 466 (7302). 29.

[17]   Javier Hernandez, Daniel McDuff, Christian Infante, Pattie Maes, Karen Quigley, and Rosalind Picard. 2016. Wearable ESM: Differences in the Experience Sampling Method Across Wearable Devices. *International Conference on Human-Computer Interaction with Mobile Devices and Services*, ACM, 195-205. https://doi.org/10.1145/2935334.2935340

[18]   Stefan E. Hormuth. 1986. The sampling of experiences in situ. *Journal of Personality*. https://doi.org/10.1111/j.1467-6494.1986.tb00395.x

[19]   S. Hosio, D. Ferreira, J. Goncalves, N. van Berkel, C. Luo, M. Ahmed, H. Flores, and V. Kostakos. 2016. Monetary Assessment of Battery Life on Smartphones. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*, 1869-1880. https://doi.org/10.1145/2858036.2858285

[20]   Gary Hsieh, Ian Li, Anind Dey, Jodi Forlizzi, and Scott E. Hudson. 2008. Using Visualizations to Increase Compliance in Experience Sampling. *Proceedings of the 10th International Conference on Ubiquitous Computing*, ACM, 164-167. https://doi.org/10.1145/1409635.1409657

[21]   Masumi Iida, Patrick E. Shrout, Jean-Philippe P. Laurenceau, and Niall Bolger. 2012. Using Diary Methods in Psychological Research. In Cooper, H. ed. *APA Handbook of Research Methods in Psychology*, American Psychological Association, 277-305.

[22]   Stephen Intille, Caitlin Haynes, Dharam Maniar, Aditya Ponnada, and Justin Manjourides. 2016. µEMA: Microinteraction-based Ecological Momentary Assessment (EMA) Using a Smartwatch. *International Joint Conference on Pervasive and Ubiquitous Computing*, ACM, 1124-1128. https://doi.org/10.1145/2971648.2971717

[23]   Max Kuhn. 2017. caret: Classification and Regression Training (R package version 6.0-78).

[24]   Reed Larson, and Mihaly Csikszentmihalyi. 1983. The Experience Sampling Method. In Csikszentmihalyi, M. ed. *Flow and the Foundations of Positive Psychology*, Wiley Jossey-Bass, 41-56.

[25]   Neal Lathia, Kiran K. Rachuri, Cecilia Mascolo, and Peter J. Rentfrow. 2013. Contextual Dissonance: Design Bias in Sensor-based Experience Sampling Methods. *International Joint Conference on Pervasive and Ubiquitous Computing*, ACM, 183-192. https://doi.org/10.1145/2493432.2493452

[26]   Inseong Lee, Jaesoo Kim, and Jinwoo Kim. 2005. Use Contexts for the Mobile Internet: A Longitudinal Study Monitoring Actual Use of Mobile Internet Services. *International Journal of Human–Computer Interaction*, 18 (3). 269-292. https://doi.org/10.1207/s15327590ijhc1803_2

[27]   Jonathan S. Lefcheck, and Robert Freckleton. 2016. piecewiseSEM: Piecewise structural equation modelling in R for ecology, evolution, and systematics. *Methods in Ecology and Evolution*, 7 (5). 573-579. https://doi.org/10.1111/2041-210X.12512

[28]   Peter Lynn. 2001. The Impact of Incentives on Response Rates to Personal Interview Surveys: Role and Perceptions of Interviewers. *International Journal of Public Opinion Research*, 13 (3). 326-336. https://doi.org/10.1093/ijpor/13.3.326

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

[30]   Abhinav Mehrotra, Jo Vermeulen, Veljko Pejovic, and Mirco Musolesi. 2015. Ask, But Don't Interrupt: The Case for Interruptibility-Aware Mobile Experience Sampling. *Adjunct Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers*, ACM, 723-732. https://doi.org/10.1145/2800835.2804397

[31]   Mulligan Casey B., Barbara Schneider, and Rustin Wolfe. 2000. Time Use and Population Representation in the Sloan Study of Adolescents *NBER Working Paper No. t0265*.

[32]   Mohamed Musthag, Andrew Raij, Deepak Ganesan, Santosh Kumar, and Saul Shiffman. 2011. Exploring Micro-Incentive Strategies for Participant Compensation in High-burden Studies. *Proceedings of the 13th International Conference on Ubiquitous Computing*, ACM, 435-444. https://doi.org/10.1145/2030112.2030170

[33]   Shinichi Nakagawa, and Holger Schielzeth. 2013. A general and simple method for obtaining R2 from generalized linear mixed-effects models. *Methods in Ecology and Evolution*, 4 (2). 133-142. https://doi.org/10.1111/j.2041-210x.2012.00261.x

[34]   Tadashi Okoshi. 2015. Attelia: Reducing user's cognitive load due to interruptive notifications on smart phones. in *Pervasive Computing and Communications (PerCom), 2015 IEEE International Conference on*, IEEE. https://doi.org/10.1109/PERCOM.2015.7146515

[35]   Veljko Pejovic, and Mirco Musolesi. 2014. InterruptMe: Designing Intelligent Prompting Mechanisms for Pervasive Applications. *International Joint Conference on Pervasive and Ubiquitous Computing*, ACM, 897-908. https://doi.org/10.1145/2632048.2632062

[36]   Martin Pielot, Bruno Cardoso, Kleomenis Katevas, Joan Serrà, Aleksandar Matic, and Nuria Oliver. 2017. Beyond Interruptibility: Predicting Opportune Moments to Engage Mobile Phone Users. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1 (3). 91:91-91:25.

[37]   Mika Raento, Antti Oulasvirta, and Nathan Eagle. 2009. Smartphones: An Emerging Tool for Social Scientists. *Sociological Methods & Research*, 37 (3). 426-454. https://doi.org/10.1177/0049124108330005

[38]   David L. Sackett. 1979. Bias in Analytic Research. *Journal of Chronic Diseases*, 32 (1). 51-63. https://doi.org/10.1016/0021-9681(79)90012-2

[39]   Clayton Shepard, Ahmad Rahmati, Chad Tossell, Lin Zhong, and Phillip Kortum. 2011. LiveLab: Measuring Wireless Networks and Smartphone Users in the Field. *ACM SIGMETRICS Performance Evaluation Review*, 38 (3). 15. https://doi.org/10.1145/1925019.1925023

[40]   Arthur Stone, Ronald Kessler, and Jennifer Haythomthwatte. 1991. Measuring Daily Events and Experiences: Decisions for the Researcher. *Journal of Personality*, 59 (3). 575-607. https://doi.org/10.1111/j.1467-6494.1991.tb00260.x

[41]   Chad Tossell, Philip Kortum, Ahmad Rahmati, Clayton Shepard, and Lin Zhong. 2012. Characterizing Web Use on Smartphones. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 2769-2778. https://doi.org/10.1145/2207676.2208676

[42]   Khai N. Truong, Julie A. Kientz, Timothy Sohn, Alyssa Rosenzweig, Amanda Fonville, and Tim Smith. 2010. The Design and Evaluation of a Task-centered Battery Interface. *Proceedings of the 12th ACM International Conference on Ubiquitous Computing*, ACM, 341-350. https://doi.org/10.1145/1864349.1864400

[43]   Niels van Berkel, Denzil Ferreira, and Vassilis Kostakos. 2017. The Experience Sampling Methods on Mobile Devices. *ACM Computing Surveys*, 50 (6). 93:91-93:40. https://doi.org/10.1145/3123988

[44]   Niels van Berkel, Jorge Goncalves, Simo Hosio, and Vassilis Kostakos. 2017. Gamification of Mobile Experience Sampling Improves Data Quality and Quantity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1 (3). 107:101-107:121. https://doi.org/10.1145/3130972

[45]   Niels van Berkel, Jorge Goncalves, Lauri Lovén, Denzil Ferreira, Simo Hosio, and Vassilis Kostakos. 2018. Effect of Experience Sampling Schedules on Response Rate and Recall Accuracy of Objective Self-Reports. *International Journal of Human-Computer Studies*, in press.

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

[47]   Rui Wang, Fanglin Chen, Zhenyu Chen, Tianxing Li, Gabriella Harari, Stefanie Tignor, Xia Zhou, Dror Ben-Zeev, and Andrew T. Campbell. 2014. StudentLife: Assessing Mental Health, Academic Performance and Behavioral Trends of College Students Using Smartphones. *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, ACM, 3-14. https://doi.org/10.1145/2632048.2632054

[48]   Jason Wiese, T. Scott Saponas, and A. J. Bernheim Brush. 2013. Phoneprioception: Enabling Mobile Phones to Infer Where They Are Kept. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 2157-2166. https://doi.org/10.1145/2470654.2481296

[49] Qiang Xu, Jeffrey Erman, Alexandre Gerber, Zhuoqing Mao, Jeffrey Pang, and Shobha Venkataraman. 2011. Identifying diverse usage behaviors of smartphone apps. *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, ACM, 329-344. https://doi.org/10.1145/2068816.2068847

[50] Yulong Yang, Gradeigh D. Clark, Janne Lindqvist, and Antti Oulasvirta. 2016. Free-Form Gesture Authentication in the Wild. *Conference on Human Factors in Computing Systems*, ACM, 3722-3735. https://doi.org/10.1145/2858036.2858270

[51] Hui Zhang, Naiji Lu, Changyong Feng, Sally W. Thurston, Yinglin Xia, and Xin M. Tu. 2011. On Fitting Generalized Linear Mixed-effects Models for Binary Responses using Different Statistical Packages. *Statistics in Medicine*, 30 (20). 2562-2572. https://doi.org/10.1002/sim.4265

[52] Xiaoyi Zhang, Laura R. Pina, and James Fogarty. 2016. Examining Unlock Journaling with Diaries and Reminders for In Situ Self-Report in Health and Wellness. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, ACM, 5658-5664. https://doi.org/10.1145/2858036.2858360

[53] Sabrina Zirkel, Julie A. Garcia, and Mary C. Murphy. 2015. Experience-Sampling Research Methods and Their Potential for Education Research. *Educational Researcher*, 44 (1). 7-16. https://doi.org/10.3102/0013189X14566879