Transactions on Computer-Human Interaction

TimeToFocus: Feedback on Interruption Durations Discourages Distractions and Shortens Interruptions

Journal:	*Transactions on Computer-Human Interaction*
Manuscript ID	TOCHI-2019-0075
Manuscript Type:	Paper
Date Submitted by the Author:	22-May-2019
Complete List of Authors:	Borghouts, Judith; University College London, UCLIC Brumby, Duncan; UCL Interaction Centre, Computer Science Cox, Anna; University College London, UCL Interaction Centre
Computing Classification Systems:	Human-centered computing, Empirical studies in HCI, Human computer interaction (HCI)

TimeToFocus: Feedback on Interruption Durations Discourages Distractions and Shortens Interruptions

Judith Borghouts[†]

UCL Interaction Centre, University College London, London, United Kingdom, judith.borghouts.14@ucl.ac.uk

Duncan P. Brumby

UCL Interaction Centre, University College London, London, United Kingdom, d.brumby@ucl.ac.uk

Anna L. Cox

UCL Interaction Centre, University College London, London, United Kingdom, anna.cox@ucl.ac.uk

[†] Corresponding author.

ABSTRACT

Many computer tasks involve looking up information from different sources. Such interruptions to a task can be disruptive. In this paper, we investigate whether giving people feedback on how long they are away from their task influences their self-interruption behaviour. We conducted a contextual inquiry on self-interruption behaviour in an office workplace. Participants were observed to postpone physical interruptions until a convenient moment in the task if they were expected to take time. In contrast, observations revealed that digital interruptions were addressed immediately; participants reported these were presumed to be quick to deal with. To increase awareness of time spent on interruptions, we developed TimeToFocus, a notification tool showing people the duration of their interruptions. A field study deployment of TimeToFocus in an office workplace found that feedback on the duration of interruptions made participants reflect on what they were doing during interruptions. They reported that they used this insight to avoid task-irrelevant activities. To confirm whether participants' perceptions of the benefit of the tool could be measured, we conducted an online experiment, where participants had to retrieve information from an email sent to their personal email addresses and enter it into a spreadsheet. Participants who used our tool made shorter interruptions, completed the spreadsheet task faster, and made fewer data entry errors. We conclude that feedback on the length of interruptions can assist users in focusing on their primary task and thus improve productivity.

CCS CONCEPTS

• **Human-centered computing → Empirical studies in HCI**.

KEYWORDS

Interruptions, workplace, notifications, contextual inquiry, field study, online experiment, focus, productivity.

1 Introduction

The fragmented nature of computer work is well-documented: people often work on several different tasks and activities throughout the day and switch between these every few minutes [13]. Though some of these switches are required to progress with work, switching away from a task can be disruptive and reduce productivity: it can slow people down, increase errors and induce stress [29]. In addition, these interruptions can trigger people to further self-interrupt their work for other off-task activities [19].

As a result, there now exists a large number of tools that aim to support people in being productive by avoiding digital distractions [25]. A common approach taken in these tools is to block distracting sources, thus removing access to the distracting material completely. Examples include tools such as FocusMe [41] and Freedom [42] which restrict access to Facebook or the internet. An interview study found that this restriction is viewed positively by people who find it difficult to self-regulate distractions [22,24,34]. However, many distracting sources such as communication tools cannot be blocked as they are needed for work. Furthermore, a field study found that blocking distractions in the workplace can also cause participants to experience higher stress, as they take fewer breaks [28].

Another common approach taken by productivity tools is to track computer usage with the aim of allowing users to perceive and reflect on their behaviour. Examples include ManicTime [43] and RescueTime [44]. Interviews revealed that it is often not clear to users what to do with reflective data [11,40]. As a result, the effectiveness of these interventions on improving focus is unclear.

Whilst recognising that interruptions can occur as part of a task, it is important to consider how long people interrupt their task. This is because longer interruptions are more disruptive than shorter ones [35], and increase the likelihood that errors are made when the task is resumed [1]. There is some evidence that feedback about people's use of time may help with task focus by reducing the duration and frequency of interruptions. For example, showing people their computer activity during the past 30 minutes reduced time spent in non-work relevant applications for office workers and students [40]. Furthermore, a message encouraging people to stay focused after an interruption reduced the number of switches to unrelated tasks during online crowdsourcing work [14]. These prior studies mostly focused on reducing task-irrelevant interruptions. What remains unknown is whether these results extend to interruptions that are required for work: for instance, many tasks involve interruptions to look up relevant information from different emails and applications.

In this paper, we look at people's switches to related activities during routine computer-based work. We specifically focus on inquiries, a type of self-interruption triggered by the need to look up task-relevant information [19]. We first conducted a contextual inquiry study (Study 1) with office workers to identify current strategies to manage interruptions. Based on the results of this study, we developed TimeToFocus, a browser notification tool which shows people the duration of their interruptions from work. We consider whether providing time feedback can help people to focus better on the task they are working on, and whether the notification can reduce the duration and number of interruptions. We evaluated the use of the tool with office workers doing data entry work (Study 2), and in an online experimental task (Study 3).

2 Background

2.1 Interruptions and Fragmentation of Work

Computer work frequently gets interrupted: on average, office workers either get interrupted or self-interrupt every three minutes [13]. People switch between tasks, but tasks themselves are also often fragmented: people have to switch between documents and applications to look up information for their task. Some interruptions can be beneficial: for example, getting the right information can have a positive impact on work [19], and short breaks can improve mood and restore energy [30]. However, frequent or longer interruptions can reduce productivity, and both controlled and in-the-wild studies have found a link between fragmented attention and a decrease in work performance [3,8].

An interview study on interruption management strategies found major differences in the level of difficulty for users to manage external versus self-interruptions. Whereas external interruptions may be ignored or deferred, self-interruptions require more self-control, and are experienced as harder to resist and as more distracting [22]. Furthermore, self-interruptions take more time to recover from than external interruptions as they can end up taking much longer than planned. When switching between computer windows, there are numerous opportunities to get distracted and get diverted from the main task. For example, when switching to communication tools, users can get tempted to answer unrelated messages as well [33]. The longer an interruption is, the more disruptive it can be, so it is important to manage time spent on interruptions from a task.

2.2 Interventions to Improve Focus

There are a number of approaches to support self-interruption management and improve people's focus. Commercial applications, such as RescueTime [44] and ManicTime [43], provide users with an overview of all their computer activities, to increase awareness of their use of time. Users can view how much time they spend on documents, websites and applications, and during which hours of the day. Little work has evaluated how effective these applications are in improving focus, and interview studies have reported a lack of engagement among users [11,40]. An interview study by Collins et al. [11] on understanding people's use of RescueTime found four barriers to explain people's lack of engagement with the data: the data lacks salience, a lack of context made it difficult to extract work patterns from the data, participants felt it was not a true representation of their actual activities, and they were not sure what actions to take based on the data.

Whittaker et al. [40] interviewed office workers and students to establish user requirements for a time awareness application, and found users were primarily interested in their current activities rather than long-term behaviour. Therefore, they developed and evaluated an application which presented users with a visualisation of the last 30 minutes of computer activity. The application reduced the time spent in email, browsing and social media, but it did not increase time spent on work and it was unclear whether it improved people's productivity. Whittaker et al. speculated that participants may already have limits on the amount of time they are prepared to spend on work, but are more flexible with the amount of time they spend on other online activities.

Other commercial tools such as Freedom and FocusMe limit access to specific sources. Kim, Cho and Lee [22] developed an intervention that allowed people to block applications and websites that they considered distracting, across devices for a fixed period. The blocking feature was viewed positively by participants who found it difficult to mitigate self-interruptions themselves. However, many distracting sources, such as web browsers and instant messaging applications, cannot be blocked during work because these need to be accessed for the current work task. To investigate how appropriate a blocking approach would be in the workplace, Mark et al. [28] conducted a field study with office workers using blocking software for one week. Participants installed software that allowed them to disable websites, and were asked to block any websites they considered distracting and nonessential to work. Several participants disliked the feeling that the software was controlling them, and rather wanted to *learn* how to gain control themselves over their work and interruptions. In line with this finding, Lyngs et al. [26] argue for supporting people to develop their own regulation strategies. They evaluated 367 apps and browser extensions that aim to support digital self-control, and found that blocking distractions or removing features were the most common approaches. Instead of preventing undesired behaviour from being triggered, Lyngs et al. identified that focusing on learning desired behaviour could be a powerful mechanism, which is largely underexplored by the reviewed self-control tools.

Other interventions suggest giving participants information during a specific task may help focus. Gould, Cox and Brumby [14] looked at people's switches to unrelated activities during an online data entry task. They found that an intervention that encouraged people to stay focused after they had self-interrupted reduced the number of switches to unrelated tasks.

While previous studies have focused on developing tools to manage task-irrelevant interruptions, little work has considered task-relevant interruptions, which are needed to progress with work and are thus difficult to avoid. For instance, many computer tasks require the user to switch between different documents, applications and computer windows. An office worker might be inputting financial information in a spreadsheet, and has to open up email to look up a relevant account number. In this scenario it can be difficult to maintain focus: people can get distracted by the need to respond to an urgent but unrelated message instead. In this paper we therefore focus on the following research question: how can people be supported in maintaining focus on work while dealing with these distracting but necessary work interruptions?

2.3 Overview of Studies

In this paper, we present three studies to explore how people can be supported in self-regulating task-required interruptions. We look at interruptions during routine data entry work, which is a common task among computer workers: information workers spend approximately 20% of their time on rote work [31]. In Study

1 we conducted a contextual inquiry to identify current interruption management strategies of office workers doing routine data entry work. Based on the findings from Study 1, we developed a browser notification tool, called TimeToFocus, that shows users how much time they spend when switching away from a task. We evaluated TimeToFocus in two settings. The aim of Study 2 was to provide a naturalistic evaluation of TimeToFocus by understanding how users would use the time information to adapt behaviour in their actual everyday work. The notification was implemented as a browser extension and deployed among nine office workers to use during data entry work. Participants could select a task they wanted to focus on. They would then receive a notification on the average switching time upon every switch away from this task. After a week of using the tool, semi-structured interviews were conducted with participants to discuss their experience of using the extension. We discuss how participants reported the use of TimeToFocus to reflect on their behaviour.

To measure the impact of TimeToFocus on interruption behaviour, in Study 3 an online experiment was conducted, of which the experimental task was framed around the type of work carried out by office workers from Study 1 and 2. Participants in Study 3 had to complete an online form by entering numeric codes. These codes had to be retrieved from a message that was sent to the participant's personal email. Participants who used TimeToFocus received notifications indicating how much time they spent away from the online form when switching to their email. The aim of Study 3 was to investigate whether the notification could reduce the duration and number of switches, and whether this then led to measurable improvements in task performance. We hypothesised that the experimental group, who received a notification, would make shorter switches than a control group, who did not receive information on how much time they spent when switching away from the primary task.

This work contributes to understanding how users can be supported to manage task-required interruptions and improve task focus. Contrary to the idea that focus can be improved by blocking distracting sources, our work shows that interruptions to distracting sources are part of the activity and are difficult to avoid. This finding has implications for the design of any tools aiming to manage work-related interruptions as well as other distractions. Our paper demonstrates how making people aware of the time spent on interruptions helps people reflect on what they were doing during an interruption, and reduce the duration of interruptions. These shorter interruptions can make people more productive by helping them to complete their work faster and more accurately.

3 Study 1: contextual inquiry

We carried out a contextual inquiry study [16] with nine office workers, a user group that often deals with interruptions and distractions [28] which can slow down work and increase errors [29]. Participants worked at four different financial administration offices based at two public universities, an environment where workers often have to interrupt their work to collect information from both digital and paper artefacts [5]. Workers in these offices deal with a lot of financial data that has to be collected from various sources and manually processed. It is important that these data items are entered accurately, but there is also time pressure to finish work on time. As a prevalent data entry task that all participants were involved in was processing expense claims, we focused on observing this task.

3.1 Method

3.1.1 Participants

Nine participants (five male) took part in the study. Ages ranged from 27 to 52 (M = 36, SD = 9); three participants wished not to disclose their age. Participants identified themselves as payroll officer, payroll and pensions' assistant, accounts assistant, research manager, and financial administrator. Their experience in their current role ranged from one to 20 years. Participants were recruited through a combination of convenience and snowball sampling. They were invited to participate via emails sent to opt-in mailing lists of Finance departments, and emails forwarded by contact persons and people who had already participated.

3.1.2 Procedure

A single session with a participant lasted approximately 2 to 2.5 hours, and participants were reimbursed £15 for their participation. All sessions were audio and video recorded. A session followed the following four stages:

1. *Interview*. Participants were briefed about the study and asked questions about the type of tasks they are involved in, the type of information sources they used, and the coping strategies to manage switching between these sources. The aim of this interview was to make the participant feel comfortable and become familiar with the study and researcher, and for the researcher to get an understanding of the participant's work and job role.
2. *Think-aloud*. In this part, the participant demonstrated processing an expense claim while thinking out loud. The participant was asked to elaborate if something interesting or unusual happened, or if the participant fell quiet.
3. *Observation*. After demonstrating the task out loud, the participant continued to process expense claims as he/she would normally without explaining what he/she was doing, while the researcher observed and took notes.
4. *Summary*. The session ended with a short interview and debriefing session. The researcher summarised findings and confirmed with the participant if these assumptions were correct. If some parts of the observation needed clarification, segments of the video recording were played back to the participant, and he/she was asked to explain what was happening during these moments.

3.1.3 Data Analysis

The audio recordings were transcribed verbatim, and handwritten notes taken during the think-aloud and observation stages were typed out. Snippets from transcripts and notes were categorised based on whether they related to the physical environment, the task sequence, the cultural work environment, work artefacts or the information flow [16]. Common types of task strategies and self-interruption strategies were grouped together and coded. Any occurrence where a user collected information for a task was considered a self-interruption strategy, as it involved temporarily leaving the main task interface to collect information. There was no pre-existing coding scheme, and codes were created based on what emerged from reading over the transcripts and notes. Video recordings were played back and used to iterate and refine the codes. Additional notes were made if anything new was observed by watching these recordings.

3.2 Findings

We made 135 observations of expense claims being processed, and identified 32 different interruption strategies to collect information. These strategies were grouped into three high-level categories: *Prepare* strategies involved collecting information before starting a data entry task, *Interrupt* strategies happened when people interrupted a data entry task to collect information, and *Postpone* strategies occurred when participants were aware they needed information, but deferred collecting it.

The interviews and observations revealed that participants maintained different interruption strategies for switching to digital versus physical information sources. To investigate these differences further, we compared observations of consulting digital information sources against observations of consulting physical information sources. The 135 observations were manually categorised for 9 participants. An observation was first coded to indicate whether it exhibited a Prepare, Interrupt, or Postpone strategy. An observation was then categorised to indicate whether it was an observation of switching to a physical or digital source. Figure 1 shows the number of observations that participants, on average, used a *Prepare*, *Interrupt*, or *Postpone* strategy for physical and digital sources during our observations.

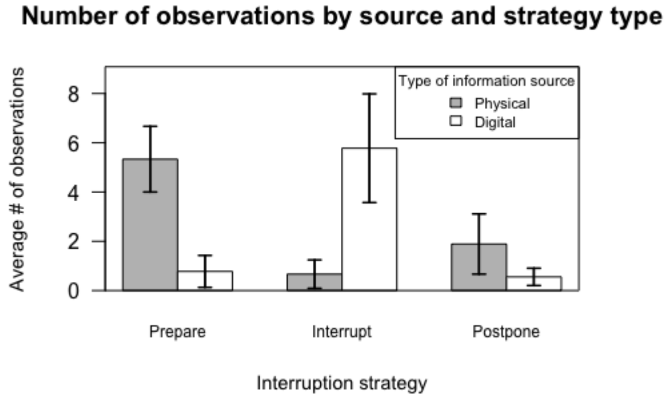

Figure 1. Bar chart showing the number of times that participants, on average, used a Prepare, Interrupt, and Postpone strategy for physical and digital information sources. In total, 135 observations were made and categorised for 9 participants. The most common strategy to collect information from physical sources was to prepare information before starting a data entry task. The most common strategy to collect information from digital sources was to interrupt and switch to the source during a data entry task.

Figure 1 reveals that strategies to collect information from physical sources were primarily grouped in either the *Prepare* or *Postpone* category, but barely in the *Interrupt* category. This means that most physical sources were prepared beforehand, or participants postponed collecting them. On the other hand, strategies to collect information from digital sources were predominantly grouped in the *Interrupt* category. This indicates that participants most often interrupted a data entry task when collecting information from digital sources.
Table 1 and 2 provide an overview of all strategies identified in the study[1]. The three columns of the tables indicate the high-level categories. Each column is filled with the observed strategies we grouped under this high-level category, and numbers in parentheses indicate for which participants we observed this behaviour. These examples are further split into rows, to indicate for which particular information source this behaviour was observed. For example, in the top row of Table 1 it can be seen that participants P1-P9 *Prepared* (column) collecting a *Paper claim form* (row) by *Placing it on their desk* (top-left cell). Each row indicates a different information source: Table 1 includes the physical sources, and Table 2 includes the digital sources. We next provide more detailed examples of some of the strategies, first for physical sources and then for digital sources.

[1] The data is available to download as a csv file at https://osf.io/u2hy9/.

Table 1. Overview of observed strategies to collect information for physical information sources. The columns indicate the three high-level categories Prepare, Interrupt and Postpone. Each column is filled with examples of observed behaviour that we categorised under this high-level category. Numbers in parentheses indicate for which participants this behaviour was observed. The rows indicate for which particular information source this behaviour was observed.

Information source	Prepare	Interrupt	Postpone
Paper claim form (9 participants)	Place on desk (P1, P2, P3, P4, P5, P6, P7, P8, P9); correct (P7); check against other sources for reliability (P1, P2, P6, P7); interpret (P3, P4); process acceptable errors (P9)	-	Send request back to claimant (P4, P9); email claimant (P5, P9); delegate to colleague (P1, P9); place note on pile on desk (P2)
Paper receipt (9 participants)	Place on desk (P1, P2, P3, P4, P5, P6, P7, P8, P9); photocopy (P4, P5, P6); check against other sources for reliability (P1); interpret (P6); annotate (P2)	-	Email claimant (P2); place note on pile on desk (P2); place in drawer (P5)
Calculator (6 participants)	Place on desk (P2, P3, P4, P5)	Retrieve from drawer (P1, P6)	-
Colleague (4 participants)	-	Ask colleague (P1, P4, P9)	Email/write note to colleague (P4, P5); delegate to colleague (P1)
Written instructions (4 participants)	Place on desk (P3, P4, P6); interpret (P3, P4, P6); check against other sources for reliability (P3, P6)	-	Email claimant (P5, P6); place in drawer (P5)
Paper personal file (2 participants)	Retrieve from shared cabinet (P1, P2)	-	Retrieve from shared cabinet (P2)
Created paper cognitive aids (2 participants)	Tape next to desk (P7)	Retrieve from drawer (P6)	-

Table 2. Overview of observed strategies to collect information for digital information sources. Each column is filled with examples of observed behaviour that we categorised under this high-level category. Numbers in parentheses indicate for which participants this behaviour was observed. The rows indicate for which particular information source this behaviour was observed.

Information source	Prepare	Interrupt	Postpone
Search engine (9 participants)	-	Look up information when needed (P1, P2, P3, P4, P5, P6, P7, P8, P9)	Stop task after not having found information (P4, P5, P7)
Spreadsheet (9 participants)	Print out document (P6, P7); create own document (P5, P6); display on second screen (P4)	Open document when needed (P1, P2, P3, P4, P8); browse (P1, P2, P4); use search option (P3); create own document (P3); memorise information (P3, P4); interleave between expenses (P4, P7, P9)	-
Currency converter application	-	Convert foreign currency (P2, P3, P4, P5, P6, P7)	-
Email inbox (6 participants)	Open email on computer (P3, P6); print out email and attachments (P3)	Look up information when needed (P1, P2, P4, P5); use search option (P1, P4); browse (P1); attend to notifications (P3, P4); read non-relevant emails (P1)	-
Intranet (5 participants)	Create own document (P4)	Look up information when needed (P2, P4, P5, P6, P7)	-
Other external websites (2 participants)	-	Look up information when needed (P4, P6)	-
Created digital cognitive aids (2 participants)	-	Look up information when needed (P4, P5)	-
PDF document (1 participant)	-	Look up information when needed (P2); browse (P2); check against other sources for reliability (P2)	Stop task after not having found information (P2)

3.2.1 Paper information sources

All participants were aware of the disruptiveness of interruptions, and the importance to focus on their data entry work. To give an example, one participant said during the interview: *'Expenses claims, (…) they do require high detail to attention. So I like to make sure that's done before I do anything else.'* (P3 - interview). As a result of the disruptiveness of interruptions, participants reported that they tried to avoid switching to unrelated tasks, and prepared most paper information sources before starting data entry work. As can be seen in the Prepare column of Table 1, participants prepared sources such as claim forms, receipts, calculators and written instructions by placing them on their desk (P1-P9), personal files were retrieved from cabinets and drawers (P1, P2), or paper sheets were already taped on walls (P7). Participants inspected these sources, and sometimes retrieved additional information sources to check the reliability, *'especially with foreign receipts, you don't really know (…) what they are.'* (P6).

A common observation was that people often discovered that they needed additional information partway through working on a data entry task. If information was nearby, for example, if it was placed in a drawer (P6) or if it could be easily handed over by a colleague (P1, P4, P9), participants interrupted their data entry task and retrieved it straight away.

If colleagues were not available and the information was situated further away, participants were observed to postpone looking for the information, and try and complete other parts of the main task first. In some cases, it was not possible to progress with the task until the required information had been found. This often stopped the task altogether, and people switched to working on a different task instead. As shown in the Postpone column of Table 1, strategies to postpone collecting information included sending the claim request back to the claimant (P4, P9), sending an email to the claimant (P2), and writing a note to a colleague who could provide the information (P4, P5). For example, *'I'm going to put this to one side. And come back to it. (…) What I do is just make a post-it note [writes post-it note], and just put it here [places it on a pile in left-hand corner of desk, and goes to new claim].'* (P2, think-aloud).

In our observations, we coded the behaviour to send an email as a 'Postpone' strategy. Participants were observed completing the rest of the task as far as possible before writing an email to get additional information. If it was not possible to continue the task without the information, they were observed stopping a task before completing it, and starting a task. We observed participants making a reminder to return to the task at a later moment, for example by placing the expense claim on a dedicated space on their desk.

3.2.2 Digital information sources

Participants tried to prepare some digital information sources beforehand as well, as illustrated at the bottom of the Prepare column of Table 2. For example, participants prepared spreadsheets by printing them out (P6, P7), displaying them on a second screen (P4), and opened a relevant email on their computer (P3, P6) in advance of starting a data entry task.

As before, people often discovered that they needed additional information partway through working on a data entry task. However, when additional information was needed from a digital information source, rather than postpone looking it up, participants were far more likely to interrupt the task and retrieve it immediately. This can be seen by looking at the bottom rows of Table 2 and comparing the Interrupt and the Postpone column: most of the strategies for digital sources are grouped under the Interrupt column, while the Postpone column is mostly empty.

Participants explained they tried to retrieve it immediately because they assumed that digital sources were easy to access and retrieving these involved little time away from the task. However, during the think-aloud sessions, as well as when discussing past incidents during the interview, it was revealed that interruptions to look up digital information could take far longer than intended, as illustrated by the following quote from P4: *"I go and make sure I've got the codes and stuff, ready to go. (…) I get halfway through and it goes, Oh, I don't know what that is. And I have to look it up. Then I'll get logged out, because it will take me longer than 5 minutes to do so."* (P4, think-aloud).

Our observations identified three main reasons for participants incurring unexpected time costs, which were corroborated by participants discussing past incidents during interviews. First, participants were observed going in and out of several documents to find what they were looking for (P1-P9), and sometimes could not find what they were looking for at all (P2, P4-P7).

Second, participants had to search through large documents with irrelevant information (e.g. spreadsheet tables with 1,000 rows and 20 columns). For example, for each expense claim, a project code had to be entered to specify for which research project the expense was made. Participants had to find this code from a large spreadsheet that contained all codes used within the organisation. During observations, participants used the search option, but also regularly did not know what specific terms to look for, and ended up scanning through the document (P2, P4, P5, P8).

Third, the irrelevant information provided potential distractions and participants were observed being diverted, for instance when they had to find information in email. Email was used by participants both as a communication tool and information source. In its role as communication tool, participants tried to ignore it during data entry work, as it was considered distracting (P1, P2, P4-P6). However, they often needed to access it to find information relevant to their work. During the think-aloud part, P1 tried to find a relevant email and opened several emails to see if it had the information he was looking for. After opening one email, he quickly knew it was not relevant but continued to read it anyway, as it reminded him of something else he had to do later on the day.

These digital interruptions had at least two negative consequences. First, the data entry system logged out after a period of inactive use, which forced participants to restart the task from the beginning: *'You'd sit down to do something, and someone (...) or something distracts you, and by the time you go back, the system's frozen and locked you out.'* (P4 - interview). Five participants reported they had experienced these logouts in the past (P4-P8), and in most cases their information was lost. Participants said the added cost of logouts kept them focused on the data entry task, and they were less likely to attend to external interruptions or switch to other, unrelated tasks. Observations however showed that each participant did interrupt their data entry and switched computer windows to look up digital information, without saving their data. Two participants were observed being logged out during the sessions (P6, P8). It was not clear to participants how long the system would wait before logging them out, or how long it would take to look up information, making it difficult to plan for these logouts: *'It doesn't time out, that's why I call it a crash out. We tend to lose various amounts of information.'* (P8, think-aloud).

A second negative consequence was that participants switched back to the wrong window, or entered the wrong information: '*If you, by mistake, left that menu, and went into another linking menu that comes up with somebody else's payroll number, you would never know that you're inputting somebody else's calculation into another record. You have to be so careful.*' (P9, think-aloud).

Though most participants had access to two screens (P3-P9) to potentially reduce window switching and mitigate its negative consequences, all window switches during data entry work happened on the same screen. Digital information was only displayed on a second screen if it was prepared beforehand and needed for a longer period of time: '*If it was a credit card claim, (...), I would have the list of credit card expenditure on one screen, and then the claim on the other. But then I'd also have another tab where I can look up codes.'* (P4, think-aloud).

3.3 Discussion

The aim of Study 1 was to investigate how people self-interrupt to access information sources for data entry work. Our results show that while people avoid task-unrelated interruptions and try to organise their data entry work so they can complete it uninterrupted, they regularly self-interrupt during the task to switch to digital information. Participants explained this behaviour by saying they expected these switches to be short, but the think-aloud sessions and the interviews revealed that interruptions often took far longer than people intended, which suggests there is a lack of awareness of time spent on digital interruptions. We first discuss possible reasons for people's behaviour, and then discuss design implications for developing a tool to manage these digital interruptions.

3.3.1 Paper versus digital interruptions

Our first finding is that participants either carefully prepared paper information sources before starting a task or postponed retrieving it but regularly interrupted themselves during the task to switch to other computer windows and find additional digital information. One possible reason for this difference in behaviour is that these switches were not experienced as 'interruptions' from the activity, but rather just another part of the *same* activity. Participants stayed on the same monitor screen when switching windows, and explained they only used a second screen for different tasks. Though participants were observed switching between computer windows to find task information, interviews revealed that they do deliberately try to minimise interruptions to unrelated tasks as data entry work requires focused attention. While this may at first seem like a contradictory finding, it is important because it provides a nuanced understanding of how people think about this type of digital interruption.

A second reason for the different management of paper versus digital interruptions is that participants had a hard time estimating the time needed for digital interruptions correctly. Prior work has shown that when participants have to choose between completing tasks with a physical or a mental effort, they choose the task that they subjectively assess as being shorter to complete, even if the objective task duration of that task is longer [37]. Potentially participants experienced the physical effort of getting up to get a paper document as more effortful or disruptive to the task, compared to the time cost of switching digital windows on the computer.

It was not always possible to prepare information, as participants did not always know they needed information until they had started a task. However, when participants realised during a task that information was needed from physical sources, they tended to postpone it. In contrast, if additional information was needed from a digital source, participants were far more likely to interrupt their task.

Window switching behaviour is consistent with previous research that has shown people switch between application windows when working on a computer every few minutes [13]. Study 1 extends these findings by making a distinction between the types of windows people switch to. Our findings suggest that even when people in this context are fairly good at reducing switches to irrelevant windows, they switch immediately to windows needed to locate information for the current task.

3.3.2 Time spent on an interruption

Another possible reason for the different treatment of paper and digital sources is the time involved in retrieving them. Participants predominantly prepared or postponed physical sources, but we observed some instances where they interrupted their work to locate information necessary to complete the task. In these cases, the information was nearby in the physical environment and retrieved rather quickly. Our findings suggest that people's decisions regarding whether or not to self-interrupt a task are influenced by the expected time involved in locating the information.

3.3.3 Distracted by other information

As described above, digital interruptions often took far longer to find than intended, as people had to spend effort finding what they needed and were distracted by other, task-irrelevant, information. A likely reason for this outcome was that people needed to access digital sources which are likely to be distracting, such as email. Arguably, participants were largely unaware of the time spent on these digital interruptions, as they adopted deferral strategies for non-digital interruptions when they perceived that they would take excessive amounts of time. This finding is important as it suggests that, even if an interruption is motivated by the goal to locate specific information and then return to the task, people can still get distracted by surrounding information. These distractions may make it difficult for people to be aware of the time that they actually spend on these interruptions – as a result, it is difficult for people to manage them effectively.

The tendency to attend to irrelevant information is similar to so-called *chains of diversion*, where the user diverts from the current task and forgets the original objective [15,17]. Previous work has explored tools that aim to prevent these diversions during a task, for example by enabling users to group windows needed for the same task [39] and disable switches to distracting sources [22]. Study 1 illustrates that these types of interventions may not be appropriate in situations where people do not know they need certain sources until they have started the task, and often need to access the sources they find distracting for work.

Study 1 provides further insight in interruptions at the workplace and how task-related interruptions, presumed to be quick and easy, can end up being time-consuming and disruptive to work. For instance, participants interrupted their work to look up digital information, even though were aware they could let logged out of the data entry system. We observed several occasions where they were logged out while trying to find information. This behaviour means we not only need to consider blocking interruptions that may be distracting from work, but also what support people can be given to control interruptions which are needed for, and considered part of, the task they want to focus on.

3.3.4 Design implications

Based on the results, we derive three main design implications for the development of our design intervention to manage task interruptions. The first implication is to increase people's awareness of their interruption behaviour. For example, upon switching windows, users may be presented with a message regarding their switching behaviour, prompting them to reflect on whether it is necessary to switch immediately, or whether they should postpone this switch to collect the required information later. Showing a message at the moment of switching windows fits with the model proposed by Lyngs et al [26], which uses the underlying cognitive mechanisms of self-regulation to frame self-regulation difficulties in ICT use. From the perspective of the model, difficulties occur because at the time of action, people's usage goals are either not strongly represented in working memory, or the value of meeting these goals is too low to control behaviour.

The second implication is that users should be in control of whether to attend to interruptions or not, as it was not always possible to avoid or defer interruptions. This is an important finding, as a common approach in many interruption management tools is to block interruptions: Lyngs et al [26] found that 40% of the 112 tools they reviewed had this functionality. Rather than blocking interruptions, workers might benefit more from tools that increase self-awareness of interruption behaviour.

The third implication is to provide people with information about the length of their digital interruptions. Participants did already effectively manage some physical interruptions, when they presumed them to be time-consuming: they addressed them before starting a task or postponed them until later. Time feedback may also help reduce time spent on interruptions, as previous studies have shown that giving users feedback on time spent on digital activities can reduce the time spent on activities away from work [14,40]. Reducing the length of interruptions can be very beneficial, as the longer people interrupt, the more disruptive it is to their main task [1].

3.4 Design intervention features

Based on the design implications, we developed the design intervention TimeToFocus to make people more aware of the time spent on interruptions. The tool was implemented as a browser extension, which allowed the user to select a main task window to focus on. Upon switching away from this window, a browser notification appeared showing users how long on average they switch away from the task window. We were interested in exploring whether this increased awareness would encourage people to reduce the number and duration of interruptions.

A first design consideration in developing the intervention was how the information should be presented to the user. Previous research found that users find it difficult to put reflective information about their use of time into context [11,40]. Instead, giving feedback about task performance during a task has been shown to make users adjust their task strategies in the moment [14,27,40]. It was therefore decided that, for the information to be most effective, it should appear during a task. To make people more aware of their interruptions, information was shown in a browser notification upon every interruption away from a primary task.

A second consideration was at which moment of an interruption the information should be shown: before or after an interruption. It was decided to show the information at the start of an interruption, so the user would be able to apply potential adjustments to their interruption behaviour immediately for that particular interruption. If information were to be shown after an interruption, the interruption would have already taken place, and the user would have to remember to change their behaviour for future interruptions.

It is not possible to provide the exact duration of an interruption that is about to take place. Therefore, the notification showed the average interruption time of all previous interruptions, to give an indication how long the interruption could take, based on past behaviour. TimeToFocus considered the average time of all digital interruptions away from a specific task, that is when the user switches from a specific task to another computer window.

TimeToFocus was evaluated in both a field and a lab setting. The notification was first evaluated in a field study in Study 2 among nine office workers, in order to investigate the subjective experience of using the tool. The aim of Study 2 was therefore to get an understanding of how people used the tool to adapt their interruption behaviour in their own everyday data entry work. We then evaluated TimeToFocus in an online experiment in Study 3 to determine its effect on productivity. The aim of Study 3 was therefore to measure whether showing people how much time they spent away from a task could reduce the duration and number of switches, and improve task performance.

4 Study 2: field study

Study 2 aimed to investigate the subjective experience, applicability and use of TimeToFocus in an office work setting. We evaluated the tool with office workers. Nine office workers were asked to install and use a browser extension implementation of TimeToFocus. The tool enabled participants to select a specific browser window they wanted to focus on. Every time they switched away from the window, the tool displayed a notification showing how long on average they were away from the window. Participants were instructed to use TimeToFocus for a week, after which they were interviewed on their experience of using the tool. The interviews aimed to explore if and how the tool could help people in managing interruptions and being more focused on their work.

4.1 Method

4.1.1 Participants

Nine participants (six female, three male) took part in the study. They were office workers at finance administration offices at one of the public universities from Study 1, and were invited to participate via emails sent to departmental mailing lists and snowballing. Participants worked in an open plan office, and seven participants occasionally worked from home. Participants' work included administrative and supportive tasks, such as processing payments, expenses, managing budgets, and responding to queries by university staff and students. The majority of participants' work was carried out in a web browser, and revolved around a number of web-based data entry systems. None of the participants had used a time or task management tool before. Participants were reimbursed with a £20 Amazon voucher after completing the study.

4.1.2 Materials

TimeToFocus was implemented as a Google Chrome extension using HTML, JavaScript and CSS. After installing the extension, an icon was permanently visible in participants' browser (see Figure 2). To use the extension, participants had to navigate to a web page in their web browser that they wanted to focus on, and click on the icon of the extension. Upon clicking on the icon, a pop-up appeared saying that the current web page was now the main task page, which indicated the start of a task session. Every time participants switched away during the session from this web page to another computer window, such as a different browser window, a document or an application, they received a notification indicating how long on average they go away for when switching away from the main task page. If participants switched away from a page for the first time, the notification showed a message that no switching data was available yet. To calculate the average switching duration, the extension recorded and saved the number and duration of switches away from the main task page for the whole session. Participants ended a session by closing the page. Due to security restrictions of browser extensions, the extension was unable to save any session data after a session had ended.

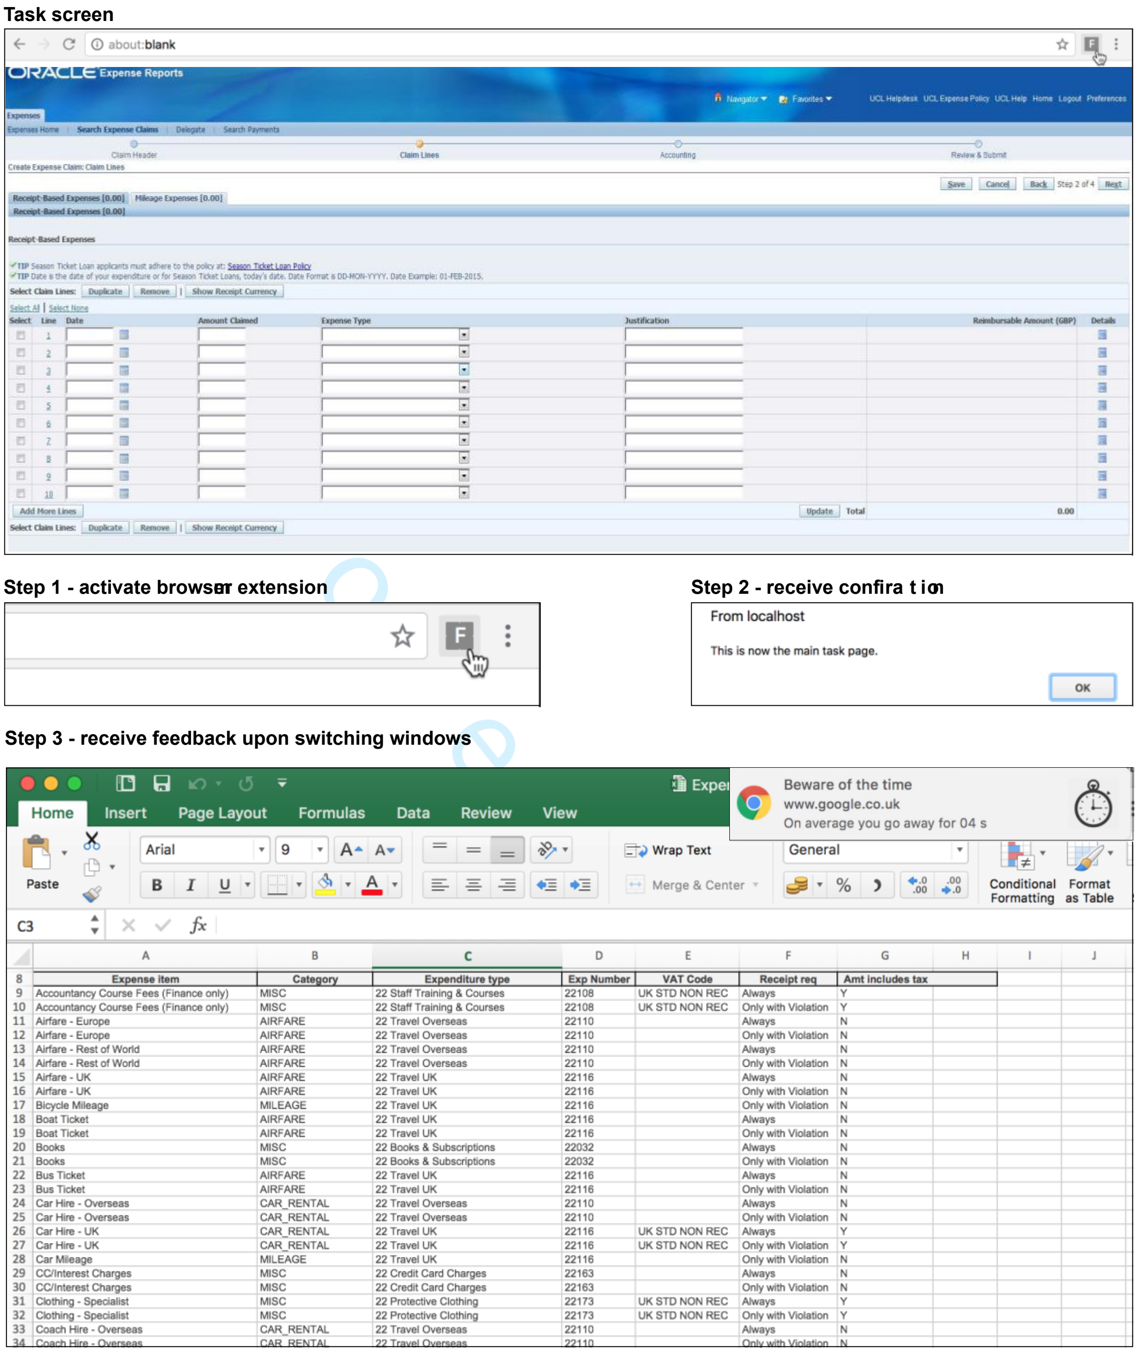

Figure 2. The browser extension implementation of TimeToFocus as shown in the browser. Participants had to navigate to a web page they wanted to focus on. They then activated the browser extension by clicking on the extension icon in their browser (Step 1). This would prompt a pop-up box to appear to confirm this web page was now selected as the main task page (Step 2). Every time the participant switched away from this page, a browser notification appeared showing how long on average they switched away from the page (Step 3).

The main focus of Study 2 was on qualitative data gathered during the interviews. We also asked participants to install ManicTime [43] on their work computer to log window switching behaviour. Given privacy concerns surrounding logging data [10], and to make sure participants were comfortable sharing their data, we made this optional. Four participants (P3, P4, P5 and P9) installed the software. A summary of logging data is included in the Results section to highlight the fragmented nature of people's work.

4.1.3 Procedure

Participants who expressed interest to take part in the study were sent an information sheet describing the full study details and a consent form to read and sign. After signing the consent form, they were sent instructions to download and install ManicTime, and an interview was scheduled after two weeks. The study was divided into two stages:

Week 1: Install ManicTime. In the first week, participants were sent instructions to install ManicTime on their work computer. Participants could pause or stop the application from running at any time. They were told that they were free to choose if, when and how often to look at the information, but that it was important to complete at least one data entry task with the application running.

Week 2: Install TimeToFocus. In the second week, participants were sent instructions to install TimeToFocus. Again, they were instructed that they were free to choose when and how often to use the browser extension, but that they had to use it for at least one data entry task.

After two weeks, participants were interviewed at either the participant's or the interviewer's office. The semi-structured interviews were structured around the following themes: how participants currently manage interruptions, tasks, time and information, the context of using TimeToFocus, the usefulness of the information provided by the extension and ManicTime, and whether they made any changes on how they managed their work. Participants were asked to share their ManicTime data. They were offered guidance and assistance on deleting or adapting any sensitive or confidential information in their data, such as application and website names. An interview lasted about 60 minutes and was audio recorded.

4.2 Findings and Discussion

Interviews were audio recorded and transcribed verbatim. Data was coded using an inductive, iterative approach of thematic analysis [7]. Analysis started after the first interview, and initial codes were refined as data collection progressed. Themes were visualised in diagrams to get insight into potential relations between different motivations and people's work practices. We first present descriptive data of people's switching behaviour as shown by the ManicTime data. We then discuss people's experiences as gathered by interview data.

4.2.1 Logged window switching behaviour

Participants' working hours differed slightly. To make the switching data comparable between participants, we only considered data between 9am and 5pm, during which all participants were at work. The mean duration of window focus was about 34 seconds, with the longest focus being 45 minutes. On average, participants made 696 computer window switches per working day. Together with the interview findings, the data shows that participants' work was characterised by short durations of focus and frequent window switches.

4.2.2 Awareness of interruption behaviour

Participants were largely aware they interrupted their work frequently and considered it the nature of their job: they regularly had to stop their work to look up task-related information and to address ad-hoc queries and requests from their department. TimeToFocus made participants realise however that they were unaware of the length of some of these interruptions. The average interruption time was much longer than they had anticipated.

Interview results suggest that common reasons for interruptions being longer than anticipated were distractions and chains of diversion [15,17], where the user further self-interrupts for other tasks. Participants tried to avoid interruptions during work that were completely unrelated, but after they had interrupted themselves for work purposes, there were opportunities to further self-interrupt for other off-task activities. The notification made people more aware of the effect this had on the duration of their interruptions: *'It's a shock, because I knew it was bad, I didn't think it was that bad. (...) So it's reflecting on, actually, a two-minute task is turning into a 15-20 minute task - why is that? (...) But again, it's distractions.'* (P9).

4.2.3 Reflecting on actions during interruptions

The increased awareness of time spent on interruptions caused people to reflect on what they were doing during past interruptions. Some interruptions were urgent, important, or necessary to progress with work, and therefore hard to avoid altogether. However, reflecting on the exact actions during the interruptions made people realise that some interruptions could be shortened, as participants often ended up getting diverted from the original goal of the interruption. For example, upon switching to their email inbox to retrieve information, participants would get diverted by reading and responding to other unread messages instead. To help remember what was happening during an interruption, P9 combined the use of TimeToFocus with the data of ManicTime: *'[The TimeToFocus notification] popped up and it said: "You go away for 7 minutes and 33 seconds. I would then have a browse [in ManicTime] And then I think: oh my gosh, I've been on emails for an hour! I haven't got anything done. So yeah, I checked it quite a lot. More so because I was so shocked. And so, I'm so interested to know, actually, what I'm doing at work.'* (P9).

Having this insight into their actions, participants tried to be more wary of potential distractions during an interruption. Sometimes the duration of an interruption was considered long, but justified. P7 was the only participant who, upon viewing the time information, was not surprised by the time she spent on work-related interruptions, and did not see any room to improve: *'To me, it doesn't kind of make me think: 'Oeh, I've been away too long'. I just think: OK, well I'm roughly aware that I've been away for an hour (…), I don't see how it kind of links with being more productive. Unless I suppose, you're really easily distracted.'* (P7).

4.2.4 Reflecting on the relevance of interruptions

Some work-related interruptions were not urgent, but participants were used to addressing them anyway if they were expected to be 'quick and easy'. These were addressed immediately so participants did not have to remind themselves to attend to it later, and it made them feel more productive if they completed more tasks. TimeToFocus made them more aware of the occurrence and actual length of these interruptions. For future interruptions, people tried to consider whether they needed to address the interruption immediately: *'I need to work on time management and (…) not spending my whole day answering irrelevant queries.'* (P9).

Participants mentioned specific sources they considered distracting, such as email, their phone, and colleagues. Some sources of distraction were not essential for work, and generally participants tried to avoid these sources during work: for example, several participants said they did not check social media at work. However, other distracting sources were used both for work and non-work purposes, such as search engines, instant messaging tools and email. It was therefore difficult to simply eliminate these distracting sources from the work environment: *'As everyone says, 'we'll just switch email off' (…). But you can bet your life that there will come a moment in whatever task you're doing you think: Oh! I have to open up email. And the moment you open up your email, that's it.'* (P2).

4.2.5 Time information for task management and perceived productivity

Completing tasks was an important component of people's work: they had an increased feeling of productivity if they explicitly ticked tasks off a list, and were driven by to-do lists and deadlines. While this could motivate people to focus on finishing a task before switching to another, it also had the contradicting effect that they interrupted their work often, if a task appeared that was considered easier to complete: *'It kind of contradicts what I told you before about (…) how I jump on them [incoming tasks] and finish them. But at the same time, it's because I don't want to have three things at once going. I want to finish, finish, finish.'* (P3).

A clear interest among participants was to not only see how much time they spent on interruptions away from the main task, but also how much time they spent on that main task overall. Currently, participants planned tasks they wanted to complete on either a daily or weekly basis, and implicitly took the time each task would take into consideration. However, people reported they often underestimated the time needed to complete tasks, a phenomenon also referred to in prior work as the planning fallacy [20]. Given the fragmented nature of participants' role and the frequency of interruptions, it was difficult to estimate how long they actually spent on tasks: *'I think that might take me 3 hours, and I'd want to get that done in one day. But yeah, obviously, things quite often take longer than I think I will, because then when I'm doing them, I might get interrupted.'* (P5).

In the same way that they used time information to reflect on whether interruptions were as long as they thought they were, they wanted to reflect on whether tasks took as long as expected. They would use this insight to be more realistic when planning tasks over time: *'Down the line, I'd think it would be extremely useful to know how much time I'm actually spending [on tasks]. Because it would help me be more productive, or be more realistic in the amount of time I need for these things to happen.'* (P3).

4.2.6 Setting goals for time limits

TimeToFocus was also used by participants to set goals on how much time they were willing to spend on interruptions. Similar to the relevance of an interruption, the appropriate length of an interruption was context-dependent as well: participants sometimes had to spend a relatively long time away from a task, for example if they had to find information in another window. Instead of absolute time limits, some participants wanted discretionary reminders to return to a task after they had reached a certain time limit, but then still have control over whether to adhere to that limit or not. Reaching the time limit could mean they were getting distracted, but it could also be the case that they were working on something relevant for work that needed more time: *'Say you have to work on that specific document, and then you end up spending half an hour on Slack chatting to your colleagues, it would be good if something's like: mate, work. Stop doing other things. But it's really hard to know what people are actually doing on these things.'* (P3).

4.2.7 Context of information

We asked participants about their use of the information provided by TimeToFocus versus ManicTime. Participants reported that the information provided by TimeToFocus was easy to read and interpret during a task. It was also clear what action to take, and participants used the information to decide whether they should reflect on past interruptions, and whether they could shorten the time away from their task. Participants looked at ManicTime at the start of the study out of curiosity and to make sure it was recording their activity correctly. However, in line with prior work [11], the extensiveness of the ManicTime data made it unclear to participants what action to take from the data, and most of them did not engage much with it for the rest of the study. It was considered too effortful and time-consuming to interpret and use the data: *'I didn't go into too much detail with it. One of the reasons is that, it would take me a lot of time and effort to use this information, to help me work better or quicker, or more efficiently. And this is either something that I don't have time to do, or I can't be bothered.'* (P3). P9 did use ManicTime, in particular to help aid her reflection on what she was doing during past interruptions, a reflection which was triggered by the notification of TimeToFocus.

Participants commented that they would have liked TimeToFocus to give additional information of their interruption behaviour over time, to place into context whether their current interruptions were longer or shorter than their usual behaviour. They would use this information to set realistic goals on interruption lengths, and to see how often they were meeting these goals.

4.2.8 Different work environments

Seven participants worked from home on occasion, and saved up tasks that required focused attention to complete at home as the office was seen as a more distracting environment. There was an implicit understanding within their department that working from home meant they needed to concentrate, and as a result, participants received fewer interruptions triggered by colleagues: *'You're working from home for a specific purpose, and therefore you don't really want to be disturbed. Unless it's absolutely urgent.'* (P2).

All participants reported there were more sources to get distracted in the office, compared to working from home. For example, participants had multiple computer screens and kept the majority of documents, browse windows and applications open on their work computer, even after they had finished with them. These windows were a further source of distraction if participants were trying to find task-related information in one of the windows: *'It's like 15 tabs, and I need to go somewhere. And I end up clicking all of them. And if there is one that is personal stuff, I end up reading it. And then five minutes after, I'm like: what was I doing? (…) So it's distracting in the way that it makes me not solely focused on one thing.'* (P3).

When working in the office, participants tried to complete tasks that required focused attention in the morning. They were more easily distracted in the afternoon, as they received more external interruptions from sources such as email, phone calls, and colleagues.

4.2.9 Measuring interruption behaviour and productivity

Overall, participants self-reported that TimeToFocus made them reflect on their interruption behaviour, and reported that they tried to avoid irrelevant activities during interruptions. While interruptions can be beneficial to productivity, interruptions that are too long can take valuable time away from work and be disruptive [1]. We therefore expect that time feedback might motivate people to shorten their interruptions, and become more focused and productive in their work. While the focus of Study 2 was on people's subjective experience of time feedback, due to privacy concerns we could not capture people's actual interruption behaviour. To evaluate whether time feedback has a measurable effect on interruption behaviour and work performance, we conducted a follow-up study where people were asked to complete an experimental task, which was designed to resemble the type of data entry work studied in Study 1 and 2.

5 Study 3: online experiment

The aim of Study 3 was to investigate whether giving people information on the duration of their window switches leads to action to reduce the number and duration of switches during a data entry task. The study used an experimental data entry task, to be able to measure whether an observed change in switching behaviour also improved task completion time and reduced data entry errors. It was however important to evaluate the notification in a setting where there was a potential to get distracted, making people more likely to be unaware of the time they spent away from a task. Designing a controlled study of self-initiated interruptions is fiendishly difficult: prior research suggests that participants tend to ignore experimenter-generated interruptions whenever they can [38]. We therefore conducted an online study, as participants get interrupted in online studies at a rate consistent with workplace observations [14]. Using this method offers the opportunity to use an experimental task in an environment that provides an abundance of natural distractions, making it a suitable method to study interruption behaviour.

A preliminary analysis of Study 3 was presented in Borghouts, Brumby and Cox [6]. Here we extend this analysis with a larger sample of participants, discuss the distribution of switching durations, analyse the effect of switching durations on task completion time, and look at inter-keystroke interval data for potential task switches outside the device.

5.1 Method

5.1.1 Participants

Forty-seven participants (30 female, 17 male) took part in the online experiment. Ages ranged from 20 to 63 (M = 29.3 years, SD = 9.1 years). The participants were recruited via university email lists, social media and online platforms to advertise academic studies, and participation was voluntary. Participants were alternately allocated to the control or experimental condition.

5.1.2 Design

The study used a between-participants design with one independent variable, a notification. In the control condition, participants did not receive a notification, but switches away from the data entry window were recorded. In the notification condition, participants were shown a notification every time they completed a trial. This notification showed how long on average they were away for when switching away from the window, before returning to the task.

The purpose of the notification was to see if the number and duration of switches could be reduced by giving participants feedback on the time spent on switches. To address this question, the dependent variables of interest were the number and duration of switches away from the data entry interface, trial completion time, and data entry errors. Switching behaviour was recorded using JavaScript's blur and focus events. These were triggered whenever a participant switched away from the data entry window, whether to their email inbox or to a different window or application.

5.1.3 Materials

The task used was based on a common routine data entry task involving processing expenses [5]. Participants were presented with an online sheet containing a set of ten 'expenses' (see Figure 3). They had to complete each row by entering the correct expense code for the expense. They retrieved this code by looking it up in a table of 25 expense categories which each had a corresponding 5-digit expense code, shown in Figure 3. Participants had to determine which category an expense belonged to, look up the code of this category and enter it in the row of the expense. We used expense categories and codes that are currently used by a public university to process expenses.

In the example in Figure 3, the expense in the top row belongs to the category 'Postage' and the participant would have to copy the code 22104 from the expense table into the empty cell of the top row. A code did not occur more than once in a trial. The codes within a trial could be entered in any order.

Once the codes of the ten expenses had been entered, participants clicked the Next button to go to the next trial and the sheet was filled with ten new expenses. In the *notification* condition, a browser notification appeared at the end of each trial at the right-hand corner of the screen that told participants the average duration of window switches away from the primary data entry task. The notification stayed visible for several seconds (a default set by the browser), or until dismissed by participants (by clicking on it).

The presentation of the notification was similar to Study 2 but differed in one important aspect. Whereas the notification in Study 2 appeared upon every switch away from the task, in Study 3 it appeared once after every trial. An early pilot study of the experiment revealed that participants switched about 10 times per trial, and experienced a notification upon every switch as too annoying. As a result, they tried to ignore the notification, diminishing the usefulness of the time feedback. The tool was therefore adapted to only appear once after every trial.

Participants were not alerted to any mistakes and once they had pressed 'Next', they could not return to the previous trial to correct any errors. Participants had to complete one practice trial, and five experimental trials. The purpose of the practice trial was for the participant to get familiar with the task, and the recorded data from this trial was excluded from the analysis.

5.1.4 Procedure

The study was advertised online with a brief description and a website link to sign up. Participants signed up for the experiment by entering their email address, and were sent an email with the table of expense categories and expense codes. The email also included instructions with a new link where the study was available. Participants were asked to complete the task on a desktop or laptop computer and open the experiment in Google Chrome, Firefox or Safari. Participants were not informed beforehand which condition they had been allocated to, and were told the purpose of the study was to understand how people perform data entry tasks. Participants in the notification condition were informed that they would receive notifications during the experiment.

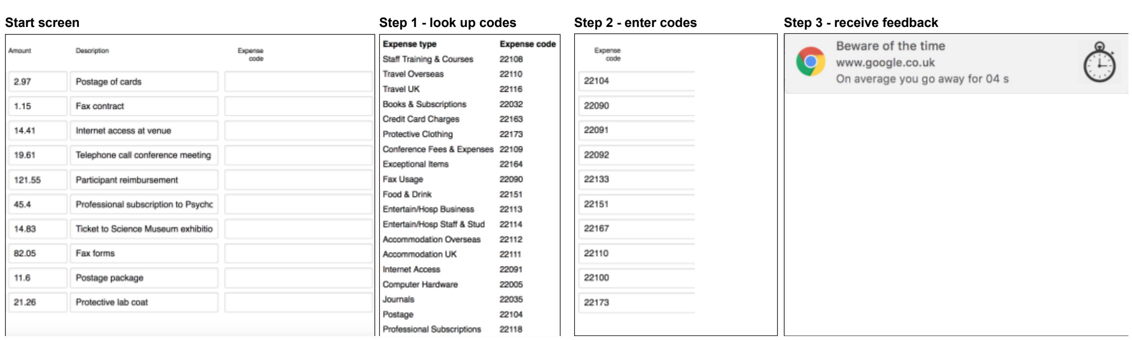

Figure 3. At the start of each trial, participants were presented with ten expenses as shown on the left. Participants had to look up codes from their email (Step 1) and enter this into a sheet (Step 2). After every trial, the notification condition received time information (Step 3).

Participants first read an online consent form on the website, and were not able to continue to the experiment until they had given explicit consent to participate. Participants in the notification condition received an additional dialog box to enable notifications in their browser, and had to click 'OK' to continue. Participants were instructed to have both their email and data entry window open on the same device, and to keep both windows maximised at all time, to ensure they had to switch back and forth between the two windows.

After completing all experimental trials, participants were shown a page of debriefing information, explaining the purpose of the study. An email address was included as a point of contact if participants had any further questions.

5.2 Findings

Participants took between 10 and 20 minutes to complete the experiment. The number of switches, length of switches and the error rate were not normally distributed, so non-parametric Mann-Whitney tests were used to analyse effects of a notification on these dependent variables. A Shapiro–Wilk test suggested that the trial completion times were normally distributed, $W = 0.94$, $p = .05$, so an independent t-test was used to analyse the effect on trial times.

5.2.1 Cleaning the data

In total, 87 participants signed up for the study. Thirty-four participants did not complete the task and their data was excluded from the dataset. Furthermore, six participants made no recorded switches and were excluded from the dataset as well. Data from the remaining 47 participants was used in the data analysis.

5.2.2 Task performance

Participants in the notification condition were faster at completing trials (M=107.61s, SD=31.15s) than participants in the control condition (M=126.27s, SD=32.61s), $t(45) = 1.98$, $p < .05$, $d = .59$. Error rates were calculated by dividing the number of data entry errors by error opportunities. The error rates were significantly lower for participants in the notification condition (M=2%, SD=2%) than participants in the control condition (M=5%, SD=5%), $U(24, 23) = 403$, $p < .01$, $r = .44$.

The notification was sent after a trial was completed, so time for participants to read and process the notification in that condition was not counted as part of the trial time. The notification appeared for at most 5s, which is considerably less than the 20s difference in trial completion times between conditions. It is therefore unlikely that reading the notification slowed participants down more than the difference in trial completion times. Nevertheless, to confirm whether people in the notification condition were actually faster in completing the overall experiment, and not just an individual trial, we also consider the total time to complete the experiment. Taking into account the total time from the start of the first trial to the end of the last trial, people in the notification condition were significantly faster (M = 8min25s, SD = 2min28s) than people in the control condition (M = 10min52s, SD = 3min13s), $U(24, 23) = 1128$, $p < .01$, $r = .63$. This analysis shows that, even though participants in the notification condition had additional time to process the information displayed by the notification, they were still significantly faster in completing the overall experiment.

5.2.3 Number and duration of switches

Participants in the notification condition made significantly shorter switches (M=4.76s, SD=1.65s) than those in the control condition (M=7.13s, SD=3.05), $U(24, 23) = 406$, $p < .01$, $r = .44$. The number of switches per trial was on average 10.6 in both conditions, and there was no significant difference in number of switches between conditions, $U(24,23)=243$, $p = .60$. As there were ten codes to be entered per trial, this suggests participants switched once for every piece of data entered.

Figure 4 shows the distribution of switching durations for each condition. For both conditions, the distribution was positively skewed with a long tail: 97% of the switches were under 20 seconds, but the longest switch was greater than seven minutes. To scale the large range of durations in one histogram, the data was plotted using a log scale for the x axis. Across conditions, there were 84 switches that were longer than 20 seconds, which highlights occurrences where participants were likely distracted.

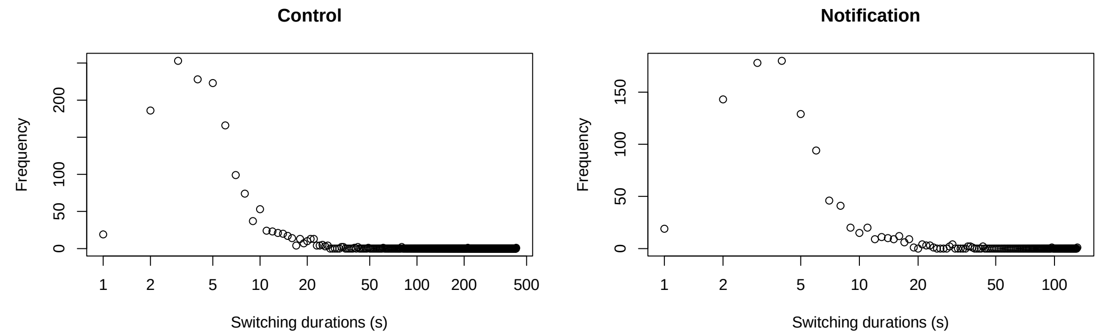

Figure 4. The distribution of switching durations in the two conditions.

5.2.4 Inter-key intervals

The primary measures to analyse switching behaviour were focus and blur events. These measures include any switch from the task window to another computer window. While this provides a good measure of window switching behaviour, it cannot capture task switches outside the device because the task window remains in focus during these task switches (e.g., a user might pause to fetch a paper document or make a cup of coffee). To help capture this broader range of instances of possible task switching behaviour, we look for longer pauses in task activity captured by an analysis of inter-keystroke interval (IKI) data. Though these intervals may have also been moments where participants had briefly paused for thought, extremely long intervals between two keystrokes may point to moments where a participant had left the task to do something else. The IKI data presented here excludes intervals where a window switch was recorded, as these moments have already been analysed in the previous section.

There was no significant difference in duration of IKIs between the notification (M = 2.02s, SD = 1.60s) and the control condition (M = 1.70s, SD = 0.91s), $U(24, 23) = 261.5$, $p = .90$. Figure 5 shows the distribution of IKIs in each condition plotted using a log scale for the x axis. The distributions were positively skewed with a long tail. The majority of IKIs were around 1-2 seconds, but there were some instances when there were long delays between keypresses: the longest measured IKI is four minutes. conditions, there were 803 instances of IKIs that were longer than five seconds. These long IKIs may have been additional task switches. However, we do not know for certain what people were doing during these instances, and what an appropriate IKI threshold would be to safely assume people had made a task switch. Therefore, we mainly focus our conclusions on our analysis of explicit window switches, and merely present the long IKIs to indicate that in addition to window switches, there may have been additional moments where people switched tasks.

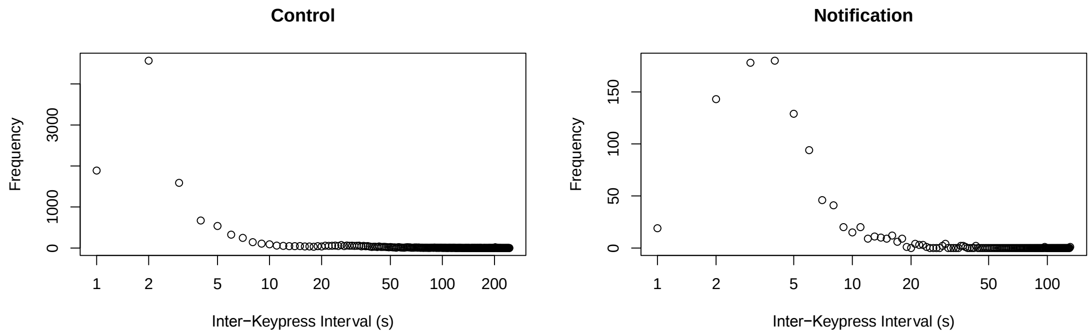

Figure 5. The distribution of inter-keypress intervals in the two conditions.

5.2.5 Effect of switching duration on task completion time

Participants who made longer switches took longer to complete trials. To see whether long switches also made people slower to resume a trial, we also consider trial completion times where the time spent on switches has been subtracted. For example, if a trial took 100 seconds, but the participant spent 20 seconds outside the task window, the adjusted trial completion time is 80 seconds.

With the adjusted trial times that has switching durations subtracted, we built a regression model to see if there was a potential relationship between trial completion times and number and duration of switches, as well as number and duration of IKIs. The linear model explained a significant amount of variation in trial times, ($R^2 = 0.41$, $F(5, 41) = 5.68$, $p < .001$). The longer people switched for, the slower they were to finish a trial, even when switching durations are subtracted from the total trial time. The number of switches explained a smaller portion of the variation, but still had a significant effect. Furthermore, the more people typed, the slower they were to complete a trial. This model provides further insight into the effect of longer switches on task performance.

5.3 Discussion

The aim of Study 3 was to see whether showing people how long they switched away from a task increased their productivity by reducing the number and length of their switches. We found that participants who received feedback on switching durations made shorter switches, were faster to complete the task, and made fewer errors. These results are important, as shortening interruptions may provide substantial benefits: it can make people less distracted and more focused on their current task, which can lead to higher productivity [18,28]. The improvements we found in data entry accuracy and completion time are also in line with previous experimental studies showing shorter interruptions improve task performance [1,35]. In this prior work, the length of the interruptions was controlled by the experimenters. Study 3 contributes to this body of work by exploring ways to shorten self-interruptions, where the length of the interruptions is controlled by users themselves.

We found that long switches significantly increased task completion time, even after subtracting the switching times. This can be explained by the resumption lag [2]: the longer people are away from a task, the longer it takes them to resume the primary task, as it takes longer to remember where they were in a task. In addition to window switches, there were also some long pauses between keypresses, which suggests that people further self-interrupted themselves outside of the computer, with the task window still in focus. However, it is difficult to know for certain what was happening during these moments in a remote study. Future studies could use additional metrics to explore what people are doing during long pauses: for example, the data entry interface may prompt the user to confirm they are still working on the task, after a certain amount of inactivity.

Most experimental studies on self-interruptions have used an artificial distraction, such as chat messages, to measure how people self-interrupt to attend to this distracting task [21,38]. Study 3 makes a methodological contribution by using participants' own personal email inbox as an interruption, based on the assumption that email provides a natural source of distraction [15,32]. Participants could complete the study at any time. Though people may receive more email during working hours, the blurring boundaries of people's working hours [23] and when people receive and manage email [9] make it difficult to control for the frequency of email across participants. Future work could examine whether time of day and frequency of incoming email affect people's self-interruption behaviour.

6 Discussion

In this paper we investigated whether showing people how long they are away from a task has an effect on people's self-interruption behaviour. Study 1 showed that physical interruptions are postponed until a convenient moment in the task if they are expected to take time, but digital interruptions are addressed immediately as these are presumed to be quick to deal with – even if these then take up longer than intended. Based on this formative study, TimeToFocus was developed showing people how long they switch away from a task. The notification tool was evaluated in two studies. Study 2 showed that feedback made people reflect on what they were doing during an interruption. They avoided interruptions that were not relevant, and set goals for how much time they were willing to spend on interruptions that were relevant. They were more focused on the intention of an interruption, and wary of potential distractions and diversions from this intention. Study 3 showed that feedback on interruption durations reduced the duration of switches, made people faster in completing a data entry task and reduced data entry errors.

In line with previous work, we found that an increased awareness of how people use time can improve focus on work [28,40]. While previous work mainly found a reduction of time spent on non-work applications, a novel finding from our work is that it can also benefit work-relevant interruptions: participants in Study 3 shortened task-related interruptions to look up relevant information.

The finding that people reflected on what they were doing during an interruption, as well as on the relevance of an interruption, is important as it may make people more in control of the intention of their interruptions. This finding builds on prior recommendations that breaks can have a positive effect on work productivity, if they are planned and not an unintended diversion [36].

Participants were also interested to see how much time they spent on a task. Prior work suggested that time information can be useful to reduce time spent on non-work activities, but not to increase time on work [40]. Rather than increasing the amount of time spent on work, we find that time information may be used to better plan *when* people spend time on tasks. It may also further show the disruptive effect of interruptions to users, as it is not only the time spent on an interruption, but also the resumption time that slows down work. This effect was also demonstrated in Study 3, where people who made longer switches were slower to complete a task, even after removing switching durations.

Participants reported they were most distracted in the office in the afternoon, which is in line with Mark et al. [31] who found that boredom in the workplace is highest in the afternoon, and people are more likely to self-interrupt when they are bored. A further likely reason we found was that there were more sources of distraction and participants were exposed to more external interruptions: prior research found that an increase in external interruptions leads to more self-interruptions [12]. On the other hand, participants felt more focused when working from home. People may therefore in particular benefit from time information at moments and settings, when they are more likely to get distracted. In Study 2, participants only used TimeToFocus at the office and it is therefore unclear whether they would use it differently when working from home. Future research could further investigate people's self-interruption behaviour in different work settings.

Our work highlights that the disruptiveness and appropriate length of an interruption not only depends on interruption properties, but also the context in which an interruption occurs. For instance, communication tools may be considered distracting and are best avoided in some situations, but in another situation the user may need to find information in email which is essential for the current task. Furthermore, though interruptions become more disruptive the longer they are, in some situations it is necessary to find information and spend a longer time away from a task. Prior approaches to block self-interruptions to distracting sources

[22,28] or impose time limits [42] are too restrictive in this situation. A more appropriate approach is to give users control over when and how long to address interruptions, and to give them useful information to help them learn how to best control and self-adjust their behaviour. Prior work found that blocking distractions at the workplace was experienced by several participants as too controlling, and participants rather wanted to learn to gain control of their work [28]. Our work makes an important contribution by showing how showing people feedback on interruption length can help in gaining control over time spent on interruptions.

6.1 Implications

Previous work has highlighted several problems with existing commercial time tracking and management applications: these often are time-consuming to use, it is not immediately clear to users what action to take based on their computer usage data, and they can restrict user activities too much by blocking distracting sources [11,40]. Our findings partly corroborate these issues, and demonstrate several pointers that can inform the design of time applications.

Presenting actionable time information. First, when providing users with a data log of their computer activities, they need to have a specific starting point of what it is they want to find out for them to be able to use it and act on it. Participants in Study 2 were not interested in their overall computer activity, but were mostly interested in the time they spent on, or away from, a specific task. By presenting a simple and precise measure, in our case the length of an interruption, participants were provided with a specific target of what to reflect on and change, and did not need to go through the effort of having to interpret information of all their activity. As some participants in Study 2 did want to have access to more detailed information about their activity during a specific interruption, a simple presentation in the moment can be complemented by a more complete log running in the background. It would also be interesting to give users control over what information they are interested in to see in the notification. For example, most participants were not only interested in the length of interruptions during a task, but also on the length of their task overall. This could help participants to better manage their tasks.

Giving time feedback during a task. Second, by showing information during the task, participants can react and change their behaviour immediately and do not have to remind themselves to look at information later [14,27]. Participants in Study 2 were prompted by the notification to reflect on what they were doing during an interruption, but often forgot to look back at their computer activities on other occasions. Participants in Study 3 were able to act on the explicit information they were given in the short time space of an experiment, which had a positive effect on their task performance.

Giving users control. Several participants wanted to set time limits on interruptions, a user need which has also been found in earlier work on interruptions [28,40]. Based on our finding however that interruptions are context-dependent, we hypothesise that imposing a strict time limit may be too restrictive. Rather, giving timed reminders to return to a task may make people more aware of the length of their interruptions, while still giving them control whether to actually return to a task or not.

Tracking behaviour over time. Lastly, a promising area to investigate would be to record the interruptions and give participants insight in how their changes have an effect over time. Although it was clear to participants in Study 2 what action they had to take based on the data presented by TimeToFocus, some felt they did not have sufficient information as to whether their actions had any effect over time.

6.2 Limitations

Due to privacy concerns, we were not able to collect logging data in Study 1 and 2 from all participants. For this reason, we conducted Study 3 to measure the impact of time feedback in the context of an online experiment. We are however unable to make any concluding claims as to whether time feedback would have any significant effect on participants' window switching and task focus behaviour over time. In addition, the presence of an observer in Study 1 may have influenced people's behaviour, and people may self-interrupt more often or for longer with no observer present. It would be useful to conduct future studies at other office settings that will allow for additional quantitative data gathering techniques.

Furthermore, TimeToFocus only focused on digital interruptions, but as was apparent in Study 1 and 2, people also deal with interruptions and distractions beyond the computer. Future work could look at also collecting and showing data from these interruptions. For example, ManicTime uses PC inactivity to indicate when the user is away from computer-based work. Other sensitive measures to detect moments where the user has likely interrupted their work could be inter-key intervals or mouse clicks.

This paper focused on routine data entry work, which is a common computer task among many office workers. We expect the results to generalise to similar types of desktop-based work, which are termed 'administrative' tasks by Bondarenko & Janssen [4]. Administrative tasks are routine tasks, of which the steps are usually the same, and are characterised by switching frequently to many different sources, but for relatively short amounts of time. As such, collecting information is considered as an interruption or subtask to support the primary task, and not regarded as a task in itself.

Lastly, though participants indicated they modified their behaviour after using TimeToFocus, it is not certain whether they based their behaviour on the specific information provided by the extension, or whether the notification simply made them reflect and become more aware of their time. Six participants in Study 2 mentioned that they compared the displayed time with their own estimated time, but we cannot separate the effect of the feedback from the notification. Controlled follow-up studies are needed to make concluding claims of causality.

7 Conclusion

Interruptions and distractions during computer-based work are common, and it can be difficult to maintain focus. However, some of these interruptions are needed to progress with work, and are difficult to avoid. In this paper, we introduced TimeToFocus to help people better manage interruptions. This work contributes to our understanding of how increased awareness of interruption durations can help users to focus on their work. Our results suggest that using TimeToFocus, participants can mitigate distractions but still keep control over their interruptions; these results can inform the design of productivity interventions to improve focus. Showing users how long they go away from a task can increase awareness of interruption behaviour, which can reduce the duration of interruptions, shorten the completion time of tasks and reduce errors.

REFERENCES

[1] Erik M. Altmann, J. Gregory Trafton, and David Z. Hambrick. 2017. Effects of Interruption Length on Procedural Errors. *J. Exp. Psychol. Appl.* 23, 2 (2017), 216–229. DOI:https://doi.org/10.1037/xap0000117

[2] Erik M. Altmann and J Gregory Trafton. 2004. Task Interruption: Resumption Lag and the Role of Cues. In *Proceedings of the 26th Annual Conference of the Cognitive Science Society*, 43–48.

[3] Brian P Bailey, Joseph A Konstan, and John V Carlis. 2001. The Effects of Interruptions on Task Performance, Annoyance, and Anxiety in the User Interface. In *INTERACT '01*, 593–601.

[4] Olha Bondarenko and R. Janssen. 2005. Documents at hand: Learning from Paper to Improve Digital Technologies. In *CHI '05*, 121–130. DOI:https://doi.org/10.1.1.137.7259

[5] Judith Borghouts, Duncan P. Brumby, and Anna L. Cox. 2017. Batching, Error Checking and Data Collecting: Understanding Data Entry in a Financial Office. In *Proceedings of 15th European Conference on Computer-Supported Cooperative Work*. DOI:https://doi.org/10.18420/ecscw2017-4

[6] Judith Borghouts, Duncan P. Brumby, and Anna L. Cox. 2018. Looking Up Information in Email: Feedback on Visit Durations Discourages Distractions. In *CHI '18 Extended Abstracts*. DOI:https://doi.org/10.1145/3170427.3188607

[7] V. Braun and V. Clarke. 2006. Using thematic analysis in psychology. *Qual. Res. Psychol.* 3, 2 (2006), 77–101. DOI:https://doi.org/10.1191/1478088706qp063oa

[8] L. Mark Carrier, Larry D Rosen, Nancy A Cheever, and Alex F Lim. 2015. Causes, effects, and practicalities of everyday multitasking. *Dev. Rev.* 35, (2015), 64–78. DOI:https://doi.org/10.1016/j.dr.2014.12.005

[9] Marta E. Cecchinato, Anna L. Cox, and Jon Bird. 2015. Working 9-5? Professional Differences in Email and Boundary Management Practices. In *CHI'15*, 3989–3998. DOI:https://doi.org/10.1145/2702123.2702537

[10] Marta E. Cecchinato, Abigail J. Sellen, M. Shokouhi, and G. Smyth. 2016. Finding Email in a Multi-Account, Multi-Device World. In *CHI '16*, 1200–1210. DOI:https://doi.org/10.1145/2858036.2858473

[11] Emily I. M. Collins, Anna L. Cox, Jon Bird, and Cassie Cornish-Tresstail. 2014. Barriers to Engagement with a Personal Informatics Productivity Tool. In *Proceedings of the 26th Australian Computer-Human Interaction Conference on Designing Futures the Future of Design (OzCHI '14)*, 370–379. DOI:https://doi.org/10.1145/2686612.2686668

[12] Laura A Dabbish, Gloria Mark, and Victor M. Gonzalez. 2011. Why Do I Keep Interrupting Myself?: Environment, Habit and Self-Interruption. In *CHI '11*, 3127–3130.

[13] Victor M. Gonzalez and Gloria Mark. 2004. "Constant, Constant, Multi-tasking Craziness": Managing Multiple Working Spheres. In *CHI '04*, 113–120.

[14] Sandy J.J. Gould, Anna L. Cox, and Duncan P. Brumby. 2016. Diminished Control in Crowdsourcing: An Investigation of Crowdworker Multitasking Behavior. *ACM Trans. Comput. Interact.* 23, 3 (2016), 19:1--19:29. DOI:https://doi.org/10.1145/2928269

[15] Benjamin V Hanrahan and Manuel A Pérez-Qu. 2015. Lost in Email: Pulling Users Down a Path of Interaction. In *CHI'15*, 3981–3984. DOI:https://doi.org/10.1145/2702123.2702351

[16] Karen Holtzblatt and Hugh Beyer. 2014. *Contextual Design: Evolved*. Morgan & Claypool Publishers.

[17] Shamsi T. Iqbal and Eric Horvitz. 2007. Disruption and Recovery of Computing Tasks: Field Study, Analysis, and Directions. In *CHI '07*, 677–686. DOI:https://doi.org/10.1145/1240624.1240730

[18] Shamsi T. Iqbal and Eric Horvitz. 2010. Notifications and Awareness: A Field Study of Alert Usage and Preferences. In *CSCW'10*, 27–30.

[19] Jing Jin and Laura A Dabbish. 2009. Self-Interruption on the Computer: A Typology of Discretionary Task Interleaving. In *CHI'09*, 1799–1808.

[20] Daniel Kahneman and Amos Tversky. 1979. Intuitive prediction: biases and corrective procedures. *TIMS Stud. Manag. Sci.* 12, (1979), 313–327.

[21] Ioanna Katidioti and Niels A Taatgen. 2013. Choice in Multitasking: How Delays in the Primary Task Turn a Rational Into an Irrational Multitasker. *Hum. Factors J. Hum. Factors Ergon. Soc.* 56, 4 (2013), 728–736. DOI:https://doi.org/10.1177/0018720813504216

[22] Jaejeung Kim, Kaist Chiwoo Cho, and Kaist Uichin Lee. 2017. Technology Supported Behavior Restriction for Mitigating Self-Interruptions in Multi-device Environments. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 3 (2017), 64:1--64:21. DOI:https://doi.org/10.1145/3130932

[23] Young-Ho Kim, Eun Kyoung Choe, Bongshin Lee, and Jinwook Seo. 2019. Understanding Personal Productivity: How Knowledge Workers Define, Evaluate, and Reflect on Their Productivity. In *CHI'19*. DOI:https://doi.org/10.1145/3290605.3300845

[24] Geza Kovacs, Drew Mylander Gregory, Zilin Ma, Zhengxuan Wu, Golrokh Emami, Jacob Ray, and Michael S Bernstein. 2019. Conservation of Procrastination: Do Productivity Interventions Save Time or Just Redistribute It? In *CHI'19*. DOI:https://doi.org/10.1145/3290605.3300560

[25] Ulrik Lyngs. 2018. A Cognitive Design Space for Supporting Self-Regulation of ICT Use. In *CHI '18 Extended Abstracts*.

[26] Ulrik Lyngs, Kai Lukoff, Petr Slovak, Reuben Binns, Adam Slack, Michael Inzlicht, Max Van Kleek, and Nigel Shadbolt. 2019. Self-Control in Cyberspace: Applying Dual Systems Theory to a Review of Digital Self-Control Tools. Chi (2019). DOI:https://doi.org/10.1145/3290605.3300361

[27] Horia A Maior, Max L. Wilson, and Sarah Sharples. 2018. Workload Alerts—Using Physiological Measures of Mental Workload to Provide Feedback During Tasks. *ACM Trans. Comput. Interact.* 25, 2 (2018), 9:1-9:30.

[28] Gloria Mark, Mary Czerwinski, and Shamsi T. Iqbal. 2018. Effects of Individual Differences in Blocking Workplace Distractions. In *CHI '18*. DOI:https://doi.org/10.1145/3173574.3173666

[29] Gloria Mark, Daniela Gudith, and Ulrich Klocke. 2008. The Cost of Interrupted Work: More Speed and Stress. In *CHI '08*, 107–110.

[30] Gloria Mark, Shamsi T. Iqbal, Mary Czerwinski, and Paul Johns. 2014. Capturing the Mood: Facebook and Face-to-Face Encounters in the Workplace. In *CSCW '14*, 1082–1094. DOI:https://doi.org/10.1145/2531602.2531673

[31] Gloria Mark, Shamsi T. Iqbal, Mary Czerwinski, and Paul Johns. 2014. Bored mondays and focused afternoons: the rhythm of attention and online activity in the workplace. In *CHI '14*, 3025–3034. DOI:https://doi.org/10.1145/2556288.2557204

[32] Gloria Mark, Shamsi T. Iqbal, Mary Czerwinski, Paul Johns, and Akane Sano. 2016. Email Duration, Batching and Self-interruption: Patterns of Email Use on Productivity and Stress. In *CHI '16*, 1717–1728.

[33] Gloria Mark, Stephen Voida, and Armand V Cardello. 2012. "A Pace Not Dictated by Electrons": An Empirical Study of Work Without Email. In *CHI '12*, 555–564.

[34] Alberto Monge Roffarello and Luigi De Russis. 2019. The Race Towards Digital Wellbeing: Issues and Opportunities. In *CHI'19*. DOI:https://doi.org/10.1145/3290605.3300616

[35] Christopher A. Monk, J. Gregory Trafton, and Deborah A. Boehm-Davis. 2008. The Effect of Interruption Duration and Demand on Resuming Suspended Goals. *J. Exp. Psychol. Appl.* 14, 4 (2008), 299–313. DOI:https://doi.org/10.1037/a0014402

[36] Alex Soojung-Kim Pang. 2016. *Rest: Why You Get More Done When You Work Less*. Basic Books.

[37] Cory Adam Potts, Stefan Pastel, and David A. Rosenbaum. 2017. How are cognitive and physical difficulty compared? *Attention, Perception, Psychophys.* 80, 2 (2017), 500–511. DOI:https://doi.org/10.3758/s13414-017-1434-2

[38] Dario D Salvucci and Peter Bogunovich. 2010. Multitasking and Monotasking: The Effects of Mental Workload on Deferred Task Interruptions. In *CHI '10*, 85–88.

[39] Greg Smith, Patrick Baudisch, George Robertson, Mary Czerwinski, Brian Meyers, Daniel Robbins, and Donna Andrews. 2003. GroupBar: The TaskBar Evolved. In *The Conference for the Computer-Human Interaction Special Interest Group of the Human Factors Society of Australia (OzCHI '03)*.

[40] Steve Whittaker, Victoria Hollis, and Andrew Guydish. 2016. "Don't Waste My Time": Use of Time Information Improves Focus. In *CHI '16*, 1729–1738.

[41] 2019. FocusMe. Retrieved March 27, 2019 from https://focusme.com/

[42] 2019. Freedom. Retrieved March 27, 2019 from https://freedom.to/

[43] 2019. ManicTime. Retrieved March 27, 2019 from https://www.manictime.com

[44] 2019. RescueTime. Retrieved March 27, 2019 from https://www.rescuetime.com