Brumby, D.P., Janssen, C.P., and Mark, G. (in press 2018) How do interruptions affect productivity? In Sadowski, C. & Zimmermann, T. (Eds.) *Rethinking Productivity in Software Engineering*. Apress.

Chapter 9

How Do Interruptions Affect Productivity?

- Duncan P. Brumby
- Christian P. Janssen
- Gloria Mark

Abstract

Work is frequently interrupted. What is known about how interruptions affect productivity? This important question has been studied using a variety of research methods, from controlled experiments designed to learn about the effect of interruptions on task performance, to analytical cognitive models that explain what makes an interruption disruptive, to in-situ observational studies that document the kinds of interruptions people experience in their actual workplaces. In this chapter, we review research on interruptions that has used these three research methods. We review what the methods entail, and what insights it has given on how interruptions affect productivity.

Introduction

When was the last time you were interrupted at work? If you use a computer for work, and if it has been more than a couple of minutes, count your blessings and be prepared for an upcoming interruption. Modern information work is punctuated by a constant stream of interruptions (González & Mark, 2004). These interruptions can be from external events (e.g., a colleague asking you a question, a message notification from a mobile device) or they can be self-initiated interruptions (e.g., going back and forth between two different computer applications to complete a task). A recent observational study of IT professionals found that some people interrupt themselves after just twenty seconds of settling into focused work (Meyer, Barton, Murphy, Zimmerman, & Fritz, 2017).

Given the omnipresence of interruptions in the modern workplace, researchers have asked what impact these have on productivity. This question has been studied in many application domains, from the hospital emergency room to the open-planned office, using a variety of different research methods.

In this chapter, we provide a brief overview of three prominent and complementary research methods that have been used to study interruptions. The methods we review are: (1) **controlled experiments** that demonstrate that interruptions take time to recover from and lead to errors, (2) **cognitive models** that offer a theoretical framework for explaining why and how interruptions are disruptive, and (3) **observational studies** that give a rich description of the kinds of interruptions that people experience in the workplace.

For each of these three research approaches, we shall explain the aim of the method, why it is relevant to the study of interruptions, and some of the key findings. Our aim is not to offer a comprehensive review of all studies in this area, but rather an introduction focusing on our own past research, which spans each of these three methods. We direct the interested reader to more comprehensive reviews of the interruptions literature (e.g., Janssen, Gould, Li, Brumby, & Cox, 2015; Salvucci & Taatgen, 2011; Trafton & Monk, 2008).

Controlled Experiments

There is a long tradition of experiments being conducted to learn about the effect of interruptions on task performance. The earliest studies were conducted back in the 1920s and focused on how well people remembered tasks that they had previously worked on. In these experiments, Zeigarnik (1927) demonstrated that people were better at recalling the details of incomplete or interrupted tasks than tasks that had been finished.

Since the advent of the computer revolution, research has focused on investigating the impact that interruptions have on task performance and productivity. This shift was probably spurred on by people's annoyances with poorly designed computer notification systems that interrupted them to attend to incoming emails or perform software updates while trying to work on other important tasks. Experiments offer a suitable research method to address the question of whether these feelings of being annoyed by interruptions and notifications translate into systematic and observable decrements in task performance.

What is the Aim of an Experiment?

Before we review what has been learnt from interruption experiments, it is worth taking a moment to reflect on the purpose of an experiment. Experiments are designed to test a hypothesis. For example, do people work slower when interrupted compared to when they have not been interrupted? To test this hypothesis, the researcher manipulates a feature of interest (the independent variable), which in our case might be the presence or absence of an interrupting task. The researcher wants to learn whether this manipulation has an effect on an outcome measure (the dependent variable), which in our case might be how quickly a task is completed.

Experiments are designed to test the *causal* relationship between variables. To do this, the researcher will attempt to control all other extraneous variables. This is why experiments are usually conducted in a controlled setting using a fixed set of instructions and tasks given to all participants who take part in the experiment. In doing so, the researcher wants to be able to isolate whether a change in the independent variable has a reliable (i.e., statistically significant) effect on the dependent variable. If an effect exists, then it should show up time and again through the independent replication of results. As we shall learn in a moment, experiments have consistently shown that interruptions negatively impact task performance.

A Typical Interruptions Experiment.

In a typical interruptions experiment, the researcher will ask a participant to work on a contrived task that they have designed. For example, the participant might be asked to use a computer interface to order some tasty donuts (Li et al., 2008). The cover story is provided to give some context to the task that the participant has been asked to work on, and it can be easily adjusted to suite the target domain of the study. For example, naval researchers have asked participants to place orders for the construction of ships (Trafton et al., 2011) and healthcare researchers have asked participants to place orders for prescription medicines (Gould, Cox, & Brumby, 2016). Regardless of the domain, the researcher gives the participant detailed instructions on how to complete the task using the interface and plenty of opportunities to practice it before starting the main part of the experiment.

In the main part of the experiment, participants will be asked to complete a number of tasks (e.g., place 10 orders for doughnuts) using the instructed procedure. While the participant is working on this task, the researcher will occasionally interrupt them and ask them to work on a secondary task instead. The secondary task might require the participant to solve some mental arithmetic problems (Li et al., 2008) or use a mouse to track a moving cursor on the screen (Monk et al., 2008). In these experiments, the arrival of this interrupting task is carefully controlled by the experimenter, and the participant is often given no choice but to switch from the primary task to the interrupting task. This is because the researcher wants to learn whether the interrupting task affects the quality and pace of the work produced on the primary task.

How is Disruptiveness of an Interruption Measured?

This discussion leads us to consider how we measure the impact of an interruption on task performance. The primary measure that has been used is the time it takes a participant to resume work on the primary task after dealing with an interruption. This time-based measure is referred to in the literature as the *resumption lag* (Altmann & Trafton, 2002; Trafton & Monk, 2008). The resumption lag measures the time it takes a person to re-engage with a task following an interruption. A longer resumption lag following an interruption reflects a general decrease in productivity: people are taking more time to complete a task, even when the time spent working on the interrupting task is deducted. In this way, the resumption lag is taken to reflect the time that is needlessly 'wasted' as a consequence of being interrupted and later having to resume an unfinished task.

Over recent years a number of experiments have been reported that use the resumption lag measure to carefully unpack which features of an interrupting task make it disruptive. Experiments have investigated whether longer interruptions are more disruptive than shorter interruptions – finding that longer interruptions result in longer resumption lags (Monk et al., 2008; Hodgetts & Jones, 2006). Studies have also been conducted to learn whether there are better or worse points in a task to be interrupted – shorter interruption lags are found when interruptions occur at natural breakpoints in a task, such as the completion of a subtask (Adamczyk & Bailey, 2004; Borst, Taatgen, & van Rijn, 2015). The content of an interrupting task also matters – interruptions that are relevant to the primary task are less disruptive than interruptions that have nothing to do with the primary task (Iqbal & Bailey, 2008; Gould, Brumby, & Cox, 2013). As we shall discuss below, the resumption lag has been explained by assuming that interruptions interfere with people's ability to remember what they were doing prior to the interruption.

Interruptions Cause Errors.

When a person resumes a task following an interruption, it often matters whether they get it right or make a mistake. Previous research has shown that interruptions increase the likelihood of errors being made on a task, in that important components of the task are either repeated or missed (Brumby, Cox, Back, & Gould, 2013; Li, Blandford, Cairns, & Young, 2008; Trafton, Altmann, & Ratwani, 2011). This finding has been taken as evidence to support the idea that following an interruption people fail to remember what they were doing in a task prior to being interrupted.

It has also been informative to consider whether there is a link between how quickly a task is resumed and the likelihood that an error is made. As discussed above, interruption researchers have generally considered a longer resumption lag to be a bad thing – reflecting time needless wasted following an interruption. In contrast, Brumby et al. (2013) found that longer resumption lags following an interruption were in fact *beneficial* in terms of reducing the occurrence of errors. This has important practical implications for the design of systems to encourage more reflective task resumption behavior in situations where interruptions are commonplace. Based on these findings, Brumby et al. developed and tested a post-interruption interface lockout that allowed users to look at the task interface but prohibited actions to be made. This interface lockout led to a significant reduction in resumption errors because it encouraged users to take the time to cognitively re-engage with a task before diving back into it and making a mistake.

Moving Controlled Experiments Out of the Lab.

A criticism that is often leveled at the kind of interruption experiments that we've reviewed above, is that the controlled setting in which they are conducted bears little resemblance to people's actual work environments and how they manage the interruptions that they experience at work. In other words, our experiments can lack *ecological validity* because an important aspect of the phenomena that we are attempting to investigate is missing. This is an important concern because it means that the results of these interruption experiments might be of limited practical value or that they might not be valid at all when taken away from the controlled setting of the lab and applied to an actual work setting.

How might an interruption experiment lack ecological validity? Interruption experiments are often conducted in controlled environments in which the researcher actively works to remove unwanted distractions and interruptions (e.g., participants will be asked to turn off their phone and give their complete attention to the researcher's task). The reason for this is that the experimenter wants to carefully control the nature and the timing of any interruptions so as to learn how they affect performance. Ironically, this desire for control presents a major threat to the ecological validity of the experiment. This is because most of the everyday interruptions that we experience are not forced but are instead *discretionary*. For example, an email notification might appear on a screen, but we can choose whether to act on it or ignore it. By using enforced interruptions, that participants have to attend to, interruption experiments can fail to capture this important aspect of the phenomena that they are attempting to study in the lab.

To overcome concerns about low ecological validity, Gould et al. (2016) has taken an approach that relaxes experimental control over the environment in which participants work in order to study how naturally occurring interruptions affect performance. To do this, Gould et al. used an online crowdsourcing platform, Amazon's Mechanical Turk, to host an interruptions experiment. Just like in a regular interruptions experiments, participants were asked to use a browser-based task interface to place orders for prescription medicines. But unlike traditional lab experiments, participants worked on this task in their regular everyday environment: an office, a coffee shop, or their home. These are naturalistic environments that are filled with everyday interruptions and distractions. In addition, workers on crowdsourcing platforms, like Amazon's Mechanical Turk, often work on multiple tasks at the same time: the environment is designed to encourage workers to complete as many tasks as possible so as to maximize their pay. This means that a competing (interrupting) task is often present, vying for the participant's attention.

By running an interruptions experiment on a crowdsourcing platform, Gould et al. (2016) found that workers switched to other tasks once every 5 minutes. This was revealed by window switch events and pauses in progression through the task. These interruptions were not inserted by the experimenter but were naturally occurring and at the discretion of the participant. Interestingly, this rate of interruptions corresponds to that seen in observational studies (e.g., González & Mark, 2004, see below for details). While these interruptions tended to be quite brief (around 30-seconds on average), Gould et al. found that they were sufficient to negatively impact performance on the primary task: participants who interrupted more often were considerably slower at completing the task, even after accounting for the time spent not working on the task. We know this only because the primary task interface was still set and under the control of the researchers; this was not a naturalistic observation study. Gould et al.'s study provides a bridge between controlled experiments and observation studies: it provides evidence that the disruptiveness of interruptions can be readily detected out in the field and that it is not an artificial product of the controlled setting used in interruption experiments.

Summary: Controlled Experiments

By conducting controlled experiments, researchers have been able to establish that task interruptions take time to recover from and lead to errors. Experiments offer an empirical approach for systematically testing whether the manipulation of an independent variable (e.g., the duration of a task interruption) has an effect on a dependent variable (e.g., the duration of the post-interruption resumption lag). Establishing whether or not the manipulation of an independent variable has an effect on the dependent variable is of both practical and theoretical value.

In practical terms, knowledge is developed about what makes an interruption disruptive, allowing practical intervention to be developed and tested. For example, Brumby et al. (2013) established that when people made faster task resumptions they were more likely to make an error. Learning about this prompted the development of an interface lockout mechanism that stopped users from resuming a task quickly following an interruption, and so reducing task errors.

In theoretical terms, experiments support the development of theories that seek to *explain* why longer interruptions result in a longer resumption lag? What is the mechanism that causes this? How can it be explained? In the next section, we turn our attention to review efforts to develop theory using cognitive models.

Cognitive Models

What are Cognitive Models?

Once findings have been made in experiments, the data and results can be used to develop theories about human behavior and thought. Cognitive models can be used to *formalize* the cumulative knowledge that is gained from experiments into *formal theories* (e.g., mathematical equations) that can *generate predictions* for future situations. For example, a mathematical model can be used to predict the likelihood that an error will be made on a task based on the duration of an interruption (e.g., Altmann & Trafton, 2002; Borst et al., 2015). Stated differently, cognitive models help to explain why and how interruptions are disruptive.

An important characteristic of cognitive models is that they generate an *exact prediction* (i.e., generate a number) as an *outcome* (e.g., likelihood of an error), given an *input* (e.g., time away from the main task), and a *formal description* of how input is *transformed* into output (i.e., a computer program that captures theory of the process of forgetting). Other more conceptual theories of interruptions (e.g., Boehm-Davis & Remington, 2009) or multitasking (e.g., Wickens, 2008) also provide insight into human behavior and thought, but typically tend to miss at least one of these three components (output, input, or transformation step), or describe them in less formal terms, such that the details that are needed to give an exact prediction are not available.

The Value (and the Devil) is in the Details…

The value of cognitive models lies in their ability to predict aspects of human behavior and thought *in detail*. Cognitive modeling aims to unravel human thought by uncovering the details and making those details open for scientific debate (Newell, 1990). As an example, take the Memory for Goals theory of forgetting (Altmann & Trafton, 2002), which has been applied to explain the results of interruption experiments. The model can be used to make a prediction for how quickly tasks will be resumed after an interruption. To do so, the model uses a mathematical function, derived from psychological theory, to determine how quickly a person will be able to recall what they were doing prior to dealing with an interruption based on the strength of this memory. The value of the model is that it gives a prediction for how quickly someone will resume a task (i.e., the resumption lag). Moreover, the general theory of memory retrieval that underpins this model helps explain *why* these resumption lags occur (namely: due to forgetting).

Since the inception of the basic Memory for Goals theory, the theory has been refined in many ways. Examples include the prediction of errors due to interruptions (Trafton, Altmann, & Ratwani, 2011), to predict performance when rapidly switching between tasks (Altmann & Gray, 2008), and to predict situations in which people need to keep the details of multiple tasks in mind (Borst et al., 2015). The initial modeling effort was crucial in this regard: by specifying a theory (of forgetting) in detail, it allowed researchers to make predictions regarding how memory impacts other settings, which could then be tested. In the end, these new experiments led to further refinements of the theory, and to an even broader understanding of the cognitive mechanisms involved in recovering from an interruption.

Although the value of cognitive models lies in the details, this is also its Achilles' heel. If a model is to be used to make predictions for a new task, then a researcher or practitioner needs to be able to specify those details ahead of time. To then specify those details, they also need to have a detailed understanding of the modeling framework and how these details should be specified within it. This is not feasible for every researcher and practitioner.

Fortunately, building on a long tradition in human-computer interaction research (Card, Moran, & Newell, 1983), more and more tools are being made to allow for predictions in applied settings, including dynamic settings such as driving (Brumby, Janssen, Kujala, & Salvucci, 2018; Salvucci, 2009). Moreover, in some cases not all details might be needed to make a prediction. For example, based on the mathematical equations behind Memory for Goals theory, recent work by Fong, Hettinger and Ratwani (2017) was able to predict the likelihood that emergency physicians resumed their original task after an interruption on their everyday emergency ward.

What Can Cognitive Models Predict About the Impact of Interruptions on Productivity?

One of the main insights to come from modeling work using the Memory for Goals theory is that the longer an interruption, the more likely it is that errors are to occur, including forgetting to resume the task altogether (and for specific cases, the models can give even more specific and exact predictions). Therefore, the implication of this work is that there is value in avoiding being interrupted.

Models can also be used to inform our understanding of discretionary self-interruptions. Previous studies have found that people often choose to interrupt themselves, switching between different activities every few minutes (Gould et al., 2016; González & Mark, 2004). For example, an information worker who is focusing on a particular work activity will still likely choose to monitor and check their email regularly, switching back and forth between application windows. How often should the person switch between these two different activities?

In our own research, we have used cognitive models to examine how the demands of a task affect the benefit of different switching strategies (i.e., how long to focus on one task before switching back to another task). We studied this in the context of a dual-task experiment in which participants had to control a dynamic task while performing a text-entry task (Farmer, Janssen, Nguyen, & Brumby, 2017; Janssen & Brumby, 2015; Janssen, Brumby, Dowell, Chater, & Howes, 2011). We used a cognitive model to identify the best possible strategy for dividing attention between these two tasks, and then compared this to what people actually choose to do in the experiments. Across several studies, we found that people were very quick at locating the best possible strategy for dividing their time between tasks. We learn from this work that people are actually pretty good at multitasking, when the relative importance of each task is made clear to them. Cognitive modeling was a vital step in this work as it was used to identify the best possible switching strategy, without this it would not have been possible to objectively benchmark how well people were multitasking.

Summary: Cognitive Models

Cognitive models develop our understanding of why and how interruptions are disruptive. They do this by instantiating theory using mathematical models and simulations. This puts into practice the ideas we have for what is causing an interruption to impact performance. Through this line of research, Memory for Goals has emerged as an important theory. The core idea is that when dealing with an interruption, people forget what it is they were working on. Resuming a task therefore involves remembering what one was doing before the interruption. By casting this as a memory retrieval process, the Memory for Goals theory is able to draw on general theories about the nature of human memory. In practical terms, cognitive models can be used to both explain existing data but also to make predictions about what will happen in novel situations or settings.

Observational Studies

The Value of Observational Studies

Whereas controlled experiments and cognitive models enable a focus on testing specific variables while controlling other factors, observational studies (also referred to as “in-situ” studies) offer ecological validity. For example, in the laboratory, the effects of interruptions may focus on a single interruption type from a single task. In a real-world environment, people generally work on multiple tasks, receiving interruptions from a range of sources. In-situ studies can serve to uncover reasons for people's behavior (i.e., the "why" of people's practices). It is a tradeoff, however, of generalizability with ecological validity. Observational studies can be very labor-intensive, limiting the scope and scale of study. Yet, with the current revolution in sensor technologies and wearables, in-situ studies are beginning to leverage these technologies for researchers to conduct observational studies at a larger scale. Nevertheless, sensors still introduce limitations on what can be observed and how the data can be interpreted.

Most in-situ studies of interruptions have been conducted in the workplace. Workplaces can be very dynamic places and interruptions can be triggered from a number of sources involving people (colleagues, phone calls, ambient conversations), and computer and smartphone notifications (e.g. email, social media, text messaging). However, interruptions can also originate from within an individual (e.g., due to mind-wandering, Mason et al, 2007).

Constant interruptions and the consequent fragmentation of work is a way of life for many information workers (Meyer et al. 2017; Czerwinski, Horvitz, & Wilhite, 2004; Mark, González, & Harris, 2005). By closely monitoring workers in-situ, it was found that people switched activities (conversations, work on computer applications, phone calls) about every three minutes on average. At a less-granular level, when activities were clustered into tasks, or "working spheres," these were found to be interrupted or switched about every 11 minutes (González & Mark, 2004). There is a relationship of length of time on task and interruptions: the longer time spent in a working sphere, the longer is the interrupting event. It has been proposed that when interruptions are used as breaks, then such longer interruptions might be due to replenishing one's mental resources (Trougakos, Beal, Green, & Weiss, 2008).

In a work environment, observations found that people self-interrupt almost as often as experiencing interruptions by an external source such as a phone call or colleague entering the office (González & Mark, 2004; Mark et al., 2005). When these field studies were done, more than a decade ago now, most self-interruptions were found to be associated with people initiating in-person interactions. Most external interruptions were also due to verbal-based interruptions from other people rather than due to notification mechanisms from their e-mail or voice mail. In more recent years, social media has become very popular in the workplace and it is very likely that the main triggers of self and external interruptions in the present-day workplace may be different.

Benefits and detriments of interruptions

Interruptions may be beneficial or detrimental. In a workplace diary study, Czerwinski et al. (2004) showed how the work context of information workers continuously changes due to interruptions. A study of corporate managers showed that while interruptions can disrupt tasks, managers appreciate the usefulness of interruptions as it provides the opportunity to get useful work-related information (Hudson, Christensen, Kellogg, & Erickson, 2002). While social media and online micro-breaks may provide numerous benefits in the workplace, field studies have shown that they create challenges due to switching contexts.

Generally, interruptions that disrupt concentration in a task, especially when they occur at a point that is not a natural breaking point for a task, can be detrimental (Iqbal, Adamczyk, Zheng, & Bailey, 2005). External interruptions cause information workers to enter into a 'chain of distraction' where stages of preparation, diversion, resumption and recovery take time away from an ongoing task (Iqbal & Horvitz, 2007). When notifications from smart phones were turned off for a week, people reported higher levels of attention (Kushlev, Proulx, & Dunn, 2016). A large cost in switching tasks on the computer is that it has been associated with higher stress (Mark, Iqbal, Czerwinski, & Johns, 2015). Yet, people are able to adjust their work practices to manage constant face-to-face interruptions, (Rouncefield, Hughes, Rodden, & Viller, 1994), as well as to manage interruptions from computer-mediated communication (Webster & Ho, 1997).

Interruptions in the workplace can also provide benefits. Longer interruptions (or work breaks), such as taking a walk in nature during work hours, have been shown to increase focus and creativity at work (Abdullah, Czerwinski, Mark, & Johns, 2016). Observational studies have identified that people use a variety of social media and news sites to take breaks to refresh and to stimulate themselves (Jin & Dabbish, 2009). However, a growing number of workplaces have policies that regulate the use of social media at work (Olmstead, Lampe, & Ellison, 2016), which can impact the ability of people to take a mental break at work.

Stress, individual differences, and interruptions.

A few field studies have examined the relationship of stress and interruptions. In a study that focused specifically on the role of email interruptions, Kushlev and Dunn (2015) found that limiting the amount of checking email significantly reduced stress. Another field study in the workplace found that cutting off email (and consequently reducing both internal and external interruptions) significantly reduced stress (Mark Voida, & Cardello, 2012). Cutting off smartphone notifications also significantly reduced inattention and symptoms of hyperactivity (Kushlev, Proulx, & Dunn, 2016). On the other hand, when email notifications were turned off, another field study showed that some individuals increased their self-interruptions to check email due to the lack of awareness of incoming emails (Iqbal & Horwitz, 2010). It is theorized that people who multitask more and who are susceptible to interruptions may have lower ability to filter out irrelevant stimuli (Carrier Rosen, Cheever, & Lim, 2015). Other individual differences have been observed, such as the personality trait of higher neuroticism with higher task switching (Mark, Iqbal, Czerwinski, Johns, & Sano, 2016).

Productivity.

Field studies suggest that higher frequency of task switching is associated with lower perceived productivity (Meyer et al., 2017; Mark et al., 2015). Several explanations have been proposed for this relationship, including the depletion of cognitive resources used in attending to interruptions, redundancy of work when reorienting back to the task (Mark et al., 2015), and that a polychronic workstyle may be contrary to what most people prefer (Bluedorn, Kaufman, & Lane, 1992).

Strategies for Dealing with Interruptions.

Observational studies reveal that people use strategies to manage interruptions. Whereas most people prefer monochronic work (finishing one task through to completion, see Bluedorn et al., 1992), the demands of the workplace result in polychronic work (i.e., the consequent switching of attention to different tasks). Because of the expectation of working in an environment with interruptions, some people have been observed to develop strategies to adapt to the unpredictability of the working environment. Participants can externalize their memory of task information, for example in the form of artifacts such as post-it notes, the email inbox (emails sent to oneself), or electronic planners, often updated throughout the day (González & Mark, 2004). The challenge with conventional electronic planners is that they are generally not designed at a level of granularity to help people recover from interruptions from a partially-completed task.

Technological solutions have also been implemented in the field to detect when people are interruptible, with the intent to minimize interruptions at inopportune times. Promising techniques tested in the field have shown that it is possible to predict when people are in cognitive states where they can be interrupted that can minimize interruptions, reduce stress and thus minimize cognitive resources needed to reorient back to a task (Züger & Fritz, 2015; Züger et al., 2017; Iqbal & Bailey, 2010; Fogarty et al, 2005).

Summary: Observational Studies

Observational studies document the kinds of interruptions that people experience in their actual workplace. These studies are resource intensive to conduct, and so often focus in on a small number of participants, giving a detailed and rich account of a particular work setting. We have learnt from observational studies that workplace interruptions are extremely commonplace. Some of these interruptions reflect the fragmented nature of work: people work on different tasks and activities through the day and this requires constant switching between them. People also seek out interactions with others – either by having conversations with colleagues or communicating through social networking sites and email. Consistent with the results from interruption experiments, observational studies also reveal that frequent interruptions result in feelings of reduced productivity. However, regular breaks from work are also necessary and people return from breaks feeling energized and ready to resume their work.

Key Insights

We have given a brief overview of three prominent and complementary research methods that have been used to study interruptions: controlled experiments, cognitive models, and observational studies. Across these three research approaches a consistent pattern of insights emerges to help us understand how interruptions affect productivity.

The key insights are:

- Interruptions can take time from which to recover from and can lead to errors.
- Shorter interruptions are less disruptive than longer interruptions.
- Interruptions delivered during a natural break in a task are less disruptive.
- Interruptions that are relevant to the current task are less disruptive.
- Resuming a task too quickly can lead to errors being made.
- All of these above characteristics of the resumption lag can be explained by an underlying memory retrieval process.
- People self-interrupt almost as often as being interrupted by external sources.
- People often work on multiple tasks at the same time and self-interruptions are important for keeping up with these different activities.
- Interruptions can cause stress, particularly email interruptions.

- Interruptions can provide an opportunity for a break to refresh, and people take longer breaks after working on a task for longer.

Key Ideas

This chapter has offered a practical and reflective account of the complementary benefits and challenges of conducting research using each of the following three methods. The main points to reflect on are that:

- Controlled experiments are designed to test a specific hypothesis, but there are challenges with designing the experiment so that it has ecological validity.
- Cognitive models offer a theoretical framework for explaining why and how things happen (e.g., how interruptions affect productivity), but these models can be complex and difficult to develop.
- Observational studies offer a rich description of situated activity, but these studies are resource intensive, and can produce an overwhelming amount of data of which to make sense.

Acknowledgements

This work was supported by the UK Engineering and Physical Sciences Research Council grants EP/G059063/1 and EP/L504889/1, by a European Commission Marie Sklodowska-Curie Fellowship H2020-MSCA-IF-2015 grant 705010, and by the U.S. National Science Foundation under grant #1704889.

References

Abdullah, S., Czerwinski, M., Mark, G., & Johns, P. (2016). Shining (blue) light on creative ability. In Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16). ACM, New York, NY, USA, 793-804. DOI: https://doi.org/10.1145/2971648.2971751

Adamczyk, P. D., & Bailey, B. P. (2004). If not now, when?: the effects of interruption at different moments within task execution. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04). ACM, New York, NY, USA, 271-278. DOI: http://dx.doi.org/10.1145/985692.985727

Altmann, E., & Gray, W. D. (2008). An integrated model of cognitive control in task switching. Psychological Review, 115, 602–639. DOI: http://doi.org/10.1037/0033-295X.115.3.602

Altmann, E., & Trafton, J. G. (2002). Memory for goals: an activation-based model. Cognitive Science, 26, 39–83. DOI: http://doi.org/10.1207/ s15516709cog2601_2

Bluedorn, A. C., Kaufman, C. F. and Lane, P. M. (1992). How many things do you like to do at once? An introduction to monochronic and polychronic time. The Executive, 6(4), 17-26. DOI: http://www.jstor.org/stable/4165091

Boehm-Davis, D. A., & Remington, R. W. (2009). Reducing the disruptive effects of interruption: a cognitive framework for analysing the costs and benefits of intervention strategies. Accident Analysis & Prevention, 41, 1124–1129. DOI: http://doi.org/10.1016/j.aap.2009.06.029

Borst, J. P., Taatgen, N. A., & van Rijn, H. (2015). What makes interruptions disruptive?: a process-model account of the effects of the problem state bottleneck on task interruption and resumption. In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15). ACM, New York, NY, USA, 2971-2980. DOI: https://doi.org/10.1145/2702123.2702156

Brumby, D. P., Janssen, C. P., Kujala, T., & Salvucci, D. D. (2018). Computational models of user multitasking. In A. Oulasvirta, P. Kristensson, X. Bi, & A. Howes (eds.) Computational Interaction Design. Oxford, UK: Oxford University Press.

Brumby, D.P., Cox, A.L., Back, J., & Gould, S.J.J. (2013). Recovering from an interruption: investigating speed-accuracy tradeoffs in task resumption strategy. Journal of Experimental Psychology: Applied, 19, 95-107. DOI: http://dx.doi.org/10.1037/a0032696

Card, S. K., Moran, T., & Newell, A. (1983). The Psychology of Human-Computer Interaction. Hillsdale, NJ: Lawrence Erlbaum Associates.

Carrier, L. M., Rosen, L. D., Cheever, N. A., & Lim, A. F. (2015). Causes, effects, and practicalities of everyday multitasking. Developmental Review, 35, 64-78. DOI: https://doi.org/10.1016/j.dr.2014.12.005

Czerwinski, M., Horvitz, E., & Wilhite, S. (2004). A diary study of task switching and interruptions. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04). ACM, New York, NY, USA, 175-182. DOI: http://dx.doi.org/10.1145/985692.985715

Farmer, G. D., Janssen, C. P., Nguyen, A. T. and Brumby, D. P. (2017). Dividing attention between tasks: testing whether explicit payoff functions elicit optimal dual-task performance. Cognitive Science. DOI: https://doi.org/10.1111/cogs.12513.

Fogarty, J., Hudson, S. E., Atkeson, C. G., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J. C., & Yang, J. (2005). Predicting human interruptibility with sensors. ACM Transactions on Computer-Human Interaction, 12, 119-146. DOI: http://dx.doi.org/10.1145/1057237.1057243

Fong, A., Hettinger, A. Z., & Ratwani, R. M. (2017). A predictive model of emergency physician task resumption following interruptions. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17). ACM, New York, NY, USA, 2405-2410. DOI: https://doi.org/10.1145/3025453.3025700

González, V. M., & Mark, G. J. (2004). "Constant, constant, multi-tasking craziness": managing multiple working spheres. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04). ACM, New York, NY, USA, 113-120. DOI: http://dx.doi.org/10.1145/985692.985707

Gould, S. J. J., Brumby, D. P., & Cox, A. L. (2013). What does it mean for an interruption to be relevant? An investigation of relevance as a memory effect. In Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 57, 149 - 153. DOI: http://dx.doi.org/10.1177/1541931213571034

Gould, S. J. J., Cox, A. L., & Brumby, D. P. (2016). Diminished control in crowdsourcing: an investigation of crowdworker multitasking behavior. ACM Transactions on Computer-Human Interaction, 23, Article 19. DOI: http://dx.doi.org/10.1145/2928269

Hodgetts, H. M., & Jones, D. M. (2006). Interruption of the Tower of London task: Support for a goal activation approach. Journal of Experimental Psychology: General, 135, 103-115. DOI: https://doi.org/10.1037/0096-3445.135.1.103

Hudson, J. M., Christensen, J., Kellogg, W. A., & Erickson, T. (2002). "I'd be overwhelmed, but it's just one more thing to do": availability and interruption in research management. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '02). ACM, New York, NY, USA, 97-104. DOI: http://dx.doi.org/10.1145/503376.503394

Iqbal, S. T., & Bailey, B. P. (2008). Effects of intelligent notification management on users and their tasks. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08). ACM, New York, NY, USA, 93-102. DOI: https://doi.org/10.1145/1357054.1357070

Iqbal, S. T., & Horvitz, E. (2007). Disruption and recovery of computing tasks: field study, analysis, and directions. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07). ACM, New York, NY, USA, 677-686. DOI: https://doi.org/10.1145/1240624.12407302007.

Iqbal, S. T., & Horvitz, E. (2010). Notifications and awareness: a field study of alert usage and preferences. In Proceedings of the 2010 ACM conference on Computer supported cooperative work (CSCW '10). ACM, New York, NY, USA, 27-30. DOI: https://doi.org/10.1145/1718918.1718926

Iqbal, S. T., Adamczyk, P. D., Zheng, X. S., & Bailey, B. P. (2005). Towards an index of opportunity: understanding changes in mental workload during task execution. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '05). ACM, New York, NY, USA, 311-320. DOI: http://dx.doi.org/10.1145/1054972.1055016

Iqbal, S.T., & Bailey, B.P. (2010). Oasis: A framework for linking notification delivery to the perceptual structure of goal-directed tasks. ACM Transactions on Computer-Human Interaction, 17, Article 15. DOI: http://dx.doi.org/10.1145/1879831.1879833

Janssen, C. P., & Brumby, D. P. (2015). Strategic adaptation to task characteristics, incentives, and individual differences in dual-tasking. PLoS ONE, 10(7), e0130009. DOI: http://doi.org/10.1371/journal.pone.0130009

Janssen, C. P., Brumby, D. P., Dowell, J., Chater, N., & Howes, A. (2011). Identifying optimum performance trade-offs using a cognitively bounded rational analysis model of discretionary task interleaving. Topics in Cognitive Science, 3, 123-139. DOI: http://dx.doi.org/10.1111/j.1756-8765.2010.01125.x

Janssen, C. P., Gould, S. J., Li, S. Y. W., Brumby, D. P., & Cox, A. L. (2015). Integrating knowledge of multitasking and Interruptions across different perspectives and research methods. International Journal of Human-Computer Studies, 79, 1-5. DOI: http://dx.doi.org/10.1016/j.ijhcs.2015.03.002

Jin, J., & Dabbish, L. (2009). Self-interruption on the computer: a typology of discretionary task interleaving. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09). ACM, New York, NY, USA, 1799-1808. DOI: https://doi.org/10.1145/1518701.1518979

Kushlev, K., & Dunn, E.W. (2015). Checking email less frequently reduces stress. Computers in Human Behavior, 43, 220-228. DOI: http://dx.doi.org/10.1016/j.chb.2014.11.005

Kushlev, K., Proulx, J., & Dunn, E.W. (2016). "Silence Your Phones": smartphone notifications increase inattention and hyperactivity symptoms. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16). ACM, New York, NY, USA, 1011-1020. DOI: https://doi.org/10.1145/2858036.2858359

Li, S. Y. W., Blandford, A., Cairns, P., & Young, R. M. (2008). The effect of interruptions on postcompletion and other procedural errors: an account based on the activation-based goal memory model. Journal of Experimental Psychology: Applied, 14, 314 –328. DOI: http://dx.doi.org/10.1037/a0014397

Mark, G., González, V., & Harris, J. (2005). No task left behind?: examining the nature of fragmented work. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '05). ACM, New York, NY, USA, 321-330. DOI: http://dx.doi.org/10.1145/1054972.1055017

Mark, G., Iqbal, S. T., Czerwinski, M., & Johns, P. (2015). Focused, aroused, but so distractible: temporal perspectives on multitasking and communications. In Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15). ACM, New York, NY, USA, 903-916. DOI: https://doi.org/10.1145/2675133.2675221

Mark, G., Iqbal, S., Czerwinski, M., Johns, P., & Sano, A. (2016). Neurotics can't focus: an in situ study of online multitasking in the workplace. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16). ACM, New York, NY, USA, 1739-1744. DOI: https://doi.org/10.1145/2858036.2858202

Mark, G., Voida, S., & Cardello, A. (2012). "A pace not dictated by electrons": an empirical study of work without email. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12). ACM, New York, NY, USA, 555-564. DOI: http://dx.doi.org/10.1145/2207676.2207754

Mason, M. F., Norton, M. I., Van Horn, J. D., Wegner, D. M., Grafton, S. T., & Macrae, C. N. (2007). Wandering minds: the default network and stimulus-independent thought. Science, 315(5810), 393-395. DOI: https://doi.org/10.1126/science.1131295

Meyer, A. N., Barton, L. E., Murphy, G. C., Zimmerman, T., & Fritz, T. (2017). The work life of developers: activities, switches and perceived productivity. IEEE Transactions on Software Engineering, 43(12), 1178 - 1193. DOI: https://doi.org/10.1109/TSE.2017.2656886

Monk, C. A., Trafton, J. G., & Boehm-Davis, D. A. (2008). The effect of interruption duration and demand on resuming suspended goals. Journal of Experimental Psychology: Applied, 14, 299-313. DOI: https://doi.org/10.1037/ a0014402

Newell, A. (1990). Unified Theories of Cognition. Cambridge, MA: Harvard University Press.

Olmstead, K., Lampe, C., & Ellison, N. (2016). Social media and the workplace. Pew Research Center. Retrieved from http://www.pewinternet.org/2016/06/22/social-media-and-the-workplace/

Rouncefield, M., Hughes, J. A, Rodden, T., & Viller, S. (1994). Working with "constant interruption": CSCW and the small office. In Proceedings of the 1994 ACM conference on Computer supported cooperative work (CSCW '94). ACM, New York, NY, USA, 275-286. DOI: http://dx.doi.org/10.1145/192844.193028

Salvucci, D. D. (2009). Rapid prototyping and evaluation of in-vehicle interfaces. Transactions on Computer-Human Interaction, 16, Article 9. DOI: http://dx.doi.org/10.1145/1534903.1534906

Salvucci, D. D., & Taatgen, N. A. (2011). The Multitasking Mind. New York, NY: Oxford University Press.

Trafton, J. G., & Monk, C. M. (2008). Task interruptions. In D. A. Boehm-Davis (Ed.), Reviews of human factors and ergonomics (Vol. 3, pp. 111–126). Santa Monica, CA: Human Factors and Ergonomics Society.

Trafton, J. G., Altmann, E. M., & Ratwani, R. M. (2011). A memory for goals model of sequence errors. Cognitive Systems Research, 12, 134–143. DOI: http://dx.doi.org/10.1016/j.cogsys.2010.07.010

Trougakos, J. P., Beal, D. J., Green, S. G., & Weiss, H. M. (2008). Making the break count: an episodic examination of recovery activities, emotional experiences, and positive affective displays. Academy of Management Journal, 51, 131-146. DOI: http://dx.doi.org/10.5465/amj.2008.30764063

Webster, J., & Ho, H. (1997). Audience engagement in multi-media presentations. SIGMIS Database 28, 63-77. DOI: https://doi.org/10.1145/264701.264706

Wickens, C. D. (2008). Multiple resources and mental workload. Human Factors, 50, 449-455. DOI: http://doi.org/10.1518/001872008X288394

Zeigarnik, B. (1927). Das Behalten erledigter und unerledigter Handlungen. Psychologische Forschung, 9, 1-85. Translated in English as: Zeigarnik, B. (1967). On finished and unfinished tasks. In W. D. Ellis (Ed.), A sourcebook of Gestalt psychology, New York: Humanities press.

Züger, M., & Fritz, T. (2015). Interruptibility of software developers and its prediction using psycho-physiological sensors. In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15). ACM, New York, NY, USA, 2981-2990. DOI: https://doi.org/10.1145/2702123.2702593

Züger, M., Corley, C., Meyer, A. N., Li, B., Fritz, T., Shepherd, D., Augustine, V., Francis, P., Kraft, N., & Snipes, W. (2017). Reducing Interruptions at Work: A Large-Scale Field Study of FlowLight. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17). ACM, New York, NY, USA, 61-72. DOI: https://doi.org/10.1145/3025453.3025662