The unexpected effect of interruptions on garden-path problems

Jennifer Sloane

Ben R. Newell

Chris Donkin

University of New South Wales

Abstract

Interruptions are a regular part of daily life. As such, most people have many rules-of-thumbs they can use to deal with the consequences of such disruption. Here, we highlight a type of environment in which the typical approaches for dealing with interruptions could fail. We focus on "garden-path problems", where the features of a decision environment cause people to apply solutions that, while applicable in other situations, are presently inappropriate. Combining results from three experiments, we show that interruptions cause people to make more errors when solving such problems, while being unaware that the interruption affected their performance. We explain why this particular effect of interruptions is specific to garden-path problems, and is not simply because interruptions cause time pressure. Finally we suggest how people might adapt the ways in which they currently treat interruptions so as to minimize their detrimental effects on garden-path problems.

Keywords: interruptions, garden-path problems, time pressure, cognitive reflection test

The unexpected effect of interruptions on garden-path problems

Introduction

Interruptions are a common occurrence in our daily lives. People in the workplace either switch tasks or get interrupted every 3 minutes (González & Mark, 2004), emergency physicians and interns get interrupted on average 7.1 times per hour (Westbrook et al., 2010), and students are distracted every 6 minutes, mostly because of social media (Rosen, Carrier, & Cheever, 2013). Interruptions, and their impacts, are multifaceted; though they have been shown to negatively affect performance (e.g. more errors, poorer quality of work) and well-being (e.g. increased stress, emotional exhaustion) (Kirmeyer, 1988; Puranik, Koopman, & Vough, 2020; Speier, Vessey, & Valacich, 2003), some interruptions have almost no effect (Borst, Taatgen, & van Rijn, 2015; Gillie & Broadbent, 1989; Sloane, Donkin, Newell, & Liang, 2019) or may actually facilitate performance (Baron, 1986; Puranik et al., 2020; Ratwani, Andrews, McCurry, Trafton, & Peterson, 2007).

Observing such a wide range of effects of interruptions should not be surprising, considering that anyone can be interrupted while doing anything, by anything, at any time. This much variability will make it difficult to explain why and how interruptions affect behavior. However, such understanding is key if we want to anticipate the impact any given interruption will have. Our aim here is to contribute to the general understanding of how and why interruptions affect behavior by exploring the consequences of interruptions in a particular kind of decision environment.

We begin by discussing some general properties of interruptions that we do understand. Interruptions matter because they divert attention away from a primary task, and thus remove task-relevant information that is active in recent memory stores. Returning to a primary task, therefore, requires the interrupted individual to restore (at least some of) that context. Studies have shown that in many situations this recreation process takes time (i.e. resumption lag) (Altmann & Trafton, 2004; Sloane et al., 2019; Trafton, Altmann, Brock, & Mintz, 2003). That is, in addition to any time spent performing the interrupting task, it takes time to retrieve the details necessary to perform the primary task (e.g., the recall of the important elements of the task, or remembering or reconstructing any plans for future actions). Thus, one aspect

common to all interruptions is that there is a time cost to resuming an interrupted, or primary, task.

This time cost associated with interruptions means that, in some contexts, we should expect interruptions to have a similar effect as that caused by time pressure. In particular, if there is a known deadline to complete the primary task, then spending time addressing the interruption ultimately leaves less time for that primary task. However, even when there is no explicit time constraint, for whatever reason, a person may choose to only allocate a set amount of time to that task (Mark, Gudith, & Klocke, 2008). As such, people will have to compensate for lost time by rushing to complete the task after being interrupted. We observe this effect in practice, for example, where clinicians spent less time performing tasks when interrupted than when uninterrupted (Westbrook et al., 2010).

Two aspects of the way that people react to time pressure are relevant here. First, people tend to make mistakes when working under time pressure (Ratcliff & Rouder, 2000; Wickelgren, 1977). Second, and especially important for our purpose, most people seem to have the general heuristic that time pressure makes them worse at tasks. To see this, consider that most people understand that being rushed is a reasonable excuse for mistakes (e.g. "Sorry, I forgot to take my keys because I was in a rush"). In experiments, we also typically find that confidence ratings are lower when collected from participants under time pressure than from those that are not (Bago & De Neys, 2019; Pleskac & Busemeyer, 2010). So, to the extent that an interruption causes time pressure (as imposed by the environment or the person) we should expect to see the same effects: people are aware they are more likely to make mistakes.

It is also generally understood that returning to a primary task may not only cost time, but may also introduce errors. Even in the absence of time pressure, the processes required to return to a primary task are prone to error. For example, a person may misremember whether a planned action was yet taken (e.g. "Did I just grab my keys?"), or they may not recall correctly certain details of a problem or a plan (B. Edwards & Gronlund, 1998; Pickrell, McDonald, Bernstein, & Loftus, 2016; Roediger & McDermott, 1995). Even when not relying on information directly recalled from memory, we might expect the comprehension of the same material to differ on subsequent reads, leading to a different understanding of the

problem than before an interruption.

What makes interruptions so interesting, and yet also challenging to study, is that in the absence of a time constraint, people can do whatever they think is necessary to avoid or fix any errors caused by an interruption. However, what the person does to recover from an interruption will depend critically on their anticipation, or awareness, of any effect that the interruption may have (or may have had). That is, to the extent that an individual is unaware of the potential problems created by an interruption, they will have no cause to take steps to fix them. Here, we highlight this general property of interruptions by showing that people are particularly susceptible to errors after interruptions for problems where a simple, but incorrect, solution is common. We will refer to such cases as garden-path, or trick, problems. This idea of garden-path problems is not novel. In fact, these types of problems have been around for decades, dating back to the 1970's and 1980's where researchers explored garden-path problems in diverse range of settings, such as physics (Clement, 1982; Cohen, 1975; Johnson & Thompson, 1981) and medical diagnosis (Johnson et al., 1981; Johnson, Moen, & Thompson, 1988).

People struggle with artificially created garden-path problems, because they are designed to create disparity between the understanding implied by their background knowledge and the correct solution. For example, consider the garden-path sentence: "the old man the boat"¹. Although this sentence is grammatically correct, its structure is designed to lead the reader to an incorrect parsing of its objects, and thus difficulty in its comprehension. In addition to the artificial garden-path problems we create for our entertainment (Bar-Hillel, Noah, & Frederick, 2018), our world also contains many situations that have a similar structure. For example, advertisers choose to promote a yogurt brand as "low-fat" to masquerade an unhealthy amount of sugar. Similarly, physicians face the challenge of potentially having to diagnose rare diseases that present with many of the same symptoms of more common ailments, and we expect police to carry out a thorough investigation even when initial evidence implies a common-place explanation.

Experimental psychologists often use the Cognitive Reflection Test (CRT) as a way of

¹ In order for this sentence to make sense, note that "the old" is the noun and "man" is the verb.

operationalising garden-path problems in experiments. For example, the famous bat and ball problem: "A bat and a ball cost \$1.10 in total. The bat costs \$1.00 more than the ball. How much does the ball cost?" (Frederick, 2005). The way this question is worded is intentionally deceiving. On an initial read, it appears to be nothing more than a simple subtraction problem ($\$1.10 - \1.00). However, by framing the price of the bat as an amount "more than" the ball, that expectation is subverted and subtraction becomes a bad solution to the problem (De Neys, Rossi, & Houdé, 2013)².

Using a modified version of the CRT, we show that interruptions cause people to be susceptible to providing incorrect answers to garden-path problems. Foreshadowing our results, we find that interruptions cause a similar drop in performance to that of time pressure. However, by looking at a combination of response patterns and confidence, we show that the deficits caused by interruptions cannot fully be explained by the addition of time pressure. While participants under time pressure systematically reduced their confidence for garden-path problems regardless of their response (i.e. correct, intuitive, or not-intuitive), we find that participants who were interrupted showed levels of confidence that were comparable to that of participants who were neither interrupted nor under time pressure.

Method

Participants in the experiments reported here were asked to complete both trick and no-trick questions while either under time pressure, or while being interrupted (or with no additional external stressor). Participants gave an answer to the questions, and indicated their confidence in their response. We report the results of overlapping conditions from three experiments. Though each experiment differed in their details, each had conditions that were equivalent for our purposes (please see the online supplementary for a complete report of the methods and results of each experiment separately). We report the details of the experiments together both for brevity, but also because the limited number of data points per participant makes us wary that any differences we observe between conditions in any one experiment may

² Note that our framing avoids the common distinction between Type 1 and Type 2 processes, we believe, without any loss of explanatory power (De Neys, 2020; Keren & Schul, 2009; Melnikoff & Bargh, 2018)

Figure 1. A visual representation of one trial. A) Reading through the word-problem with the ability to navigate back and forth between pages in each of the three conditions. B) Selecting a response. C) Rating one's confidence.

be due to sampling variability rather than a reliable effect.

Participants

Participants were recruited through Amazon's Mechanical Turk platform, where they were paid \$2.50 for their time. Each experiment took roughly 10 minutes. Only English speakers from the United States could participate in this study. 152 participants were recruited for Experiment 1, 309 for Experiment 2, and 194 for Experiment 3. 32 participants were excluded from all analyses for reasons such as failing the attention check or not completing the entire task. This left a combined sample size of 623 participants across the three experiments.

Materials

We used 10 more recently developed and validated CRT items from Primi, Morsanyi, Chiesi, Donati, and Hamilton (2016) and Thomson and Oppenheimer (2016) as our stimuli (see Table 3 in the Appendix for the stimuli used). Five of the questions were trick questions, which, like the original CRT items, are designed to provoke a garden-path, or "intuitive", incorrect answer. The remaining four no-trick questions were word problems with no intuitive, incorrect response. We also included a final question that served as an attention check to ensure participants were engaged with the task, but was removed for all analyses.

Procedure

All participants began the experiments by filling out an online consent form. After reading the instructions, everyone had an opportunity to complete one or two practice problems to get familiar with the experimental paradigm. Then, participants completed a series of multiple-choice word problems that were a combination of trick and no-trick questions. For Experiment 1, all participants completed all four no-trick questions and all five trick questions. From Experiments 2 and 3, we analyse the data from the first two trick and two no-trick questions that participants completed (we exclude the data from questions answered after that point, because participants in those experiments received feedback on their performance after the first four questions).

Rather than having each word problem presented on one page, we split up every question to span across a total of three pages, as shown in Figure 1A. Participants first had the option to navigate between the pages as many times as they pleased by clicking the "Next" and "Previous" buttons on the screen. After they had read the three question pages, they pressed Next to navigate to the response stage as shown in Figure 1B. Participants were informed that once at the question page, they could no longer navigate to previous pages, and that they should select a response from the three available options. Finally, they were asked to rate their confidence on a slider scale whose range was described as being from "uncertain/guessing" to "absolutely certain" (see Figure 1C). Participants were then given the option to take a short break before proceeding onto the next question.

Participants were randomly assigned to one of three conditions: baseline, interruption, or time pressure (note that no participants in Experiment 3 completed the time pressure condition). Participants in the baseline condition completed the study as described above. In Experiment 1, we collected data for the baseline condition first as we needed information to finalize the design for the other two conditions. In the time pressure condition, there was a countdown timer at the top of every page as participants read the word problems and selected responses. Because the word problems were different lengths, we needed to find the average time participants spent reading each question in order to apply an appropriate countdown timer for every stimulus. To do this, we calculated how long participants spent reading all

three pages in the baseline condition for every question. Then, we took the time at the 50th percentile and rounded to the nearest second to create the timers for each word problem. The time limit for each stimulus is reported in the supplementary material.

We again used results from the baseline condition to set exactly when the interruptions would occur. On one hand, we did not want to interrupt participants too early because they may not have had enough time to read crucial pieces of information. On the other hand, we also did not want to interrupt them too late because then it would be likely that people would complete the problem before the interruption happened and would, therefore, miss the interruption altogether. Therefore, we found the average time it took 75% of participants to finish reading the first page. Then, we rounded to the nearest second and set the interruptions to happen at each of those respective times (see supplementary material). This ensured that the majority of participants should at least be on the second page when they got interrupted.

When participants were interrupted, a pop-up box appeared on their screen where they had to enter the answer to a math problem using the keyboard. Because the word problems generally required some arithmetic, we decided to use math problems as the interrupting task so that both tasks required similar cognitive processes, which would likely make the interruption more disruptive (Gillie & Broadbent, 1989; Wickens, 1991). The math problems were always simple addition problems, but were randomly generated for each interruption. Each addition problem consisted of one number between one and ten and one number between one and 100. During the interruption, participants were unable to proceed with the word problem. However, once participants responded to the math question, the pop-up box disappeared and they resumed reading the word problem exactly where they left off.

Analysis

We excluded 12% of trials on which participants in the interruption condition responded quickly enough not to be interrupted. Using the data from each participant in all three experiments, we model whether responses to each question were correct or not, X , as Bernoulli distributed,

$$X \sim \text{Bern}(\theta_{ij})$$

The probability of a correct response, θ_{ij} , is the inverse-logit function of Z_{ij} , and is estimated separately for each item, i , and for each of the baseline, interruption, and time pressure conditions, j . We use Standard Normal distributions as the priors for each Z parameter. We focus our analysis on the θ_{ij} parameters, averaged separately over trick and no-trick questions.

Again, across the three experiments, participants' confidence ratings for each question, C , are scaled to be between 0.001 and 0.999, and then modelled as Beta distributed

$$C \sim \text{Beta}(\alpha_{ij}, \beta_{ij})$$

We use a Beta distribution because it allows us to accommodate the extreme skew of the distribution of confidence ratings across participants and items. We estimated α_{ij} and β_{ij} parameters separately for each item, i , and for each of the baseline, interruption, and time pressure conditions, j . The priors placed on α and β parameters are Uniform distributions from 0 to 500. We focus our analysis on the mean of the fitted beta distribution, $\gamma = \frac{\alpha}{\alpha + \beta}$. Like θ , we average γ_{ij} parameters separately over trick and no-trick questions. Note that model parameter estimates shown here were created using JAGS, drawing 3000 samples in each of 5 chains (after 1000 samples of burn in), with no convergence or autocorrelation issues observed.

To compare γ and θ parameters across baseline, interruption, and time pressure conditions, we examine the posterior distribution of their differences. In particular, for pairs of conditions, we create a posterior distribution of the difference in a given parameter between the two conditions, and ask how many of the samples drawn from that posterior fall on one side of zero. To the extent that the difference between the two conditions is large, then we expect a small proportion of the samples drawn from the posterior distribution to be less than zero. On the other hand, if the posterior distributions from each condition are largely overlapping, then we expect roughly half of the samples from the posterior distribution of their difference to be on either side of zero. We report this p_{samples} as an index of the difference between conditions, with values closer to zero reflecting a larger difference.

Results

Trick Questions

Focusing first on trick questions, looking at the top-left panel of Figure 2, we see that participants made more correct responses in the baseline condition (72%) than in the interruption and time pressure conditions (65% and 61%, respectively). The estimates of θ from our model, averaging over the trick question items, are plotted in black. The error bars show the central 95% of the posterior distribution, and so the extent to which the error bars do not overlap suggests the differences we observe are large, relative to our measure of error. Indeed, the difference between baseline and interruption, and baseline and time pressure conditions both yielded $p_{samples}$ values of less than 0.0001. On the other hand, accuracy in interruption and time pressure conditions was much more similar ($p_{samples} = 0.20$).

While both interruption and time pressure led to worse performance for trick questions, the way the manipulations affected responses was different. The proportion of correct, intuitive, and other (not-intuitive but incorrect) responses are also plotted in this figure as the stacked grey bars, with exact values reported in Table 1. Both the interruption and time pressure manipulations led to more intuitive, incorrect responses (29% and 28%, respectively) than the baseline condition (23%). However, only the time pressure caused an increase in not-intuitive errors (12%) compared to baseline and interruption conditions (5% and 6%, respectively).

The confidence ratings reported in the bottom-left panel of Figure 2 also suggest that interruptions and time pressure had different effects on performance. We see that average confidence ratings for those in the interruption condition ($\gamma = 0.83$) were quite similar to the baseline condition ($\gamma = 0.85$), especially relative to the time pressure condition ($\gamma = 0.72$). Indeed, the posterior distribution of the difference between baseline and interruption conditions is more similar ($p_{samples} = 0.04$) than that between baseline and time pressure conditions ($p_{samples} < 0.0001$). The grey bar plots in the bottom-left panel of Figure 2 show the mean confidence ratings for each type of response. We see that confidence ratings for baseline and interruption conditions are relatively similar, while those under time pressure condition reduced their confidence for all responses.

	Trick Questions			No-Trick Questions	
	Correct	Intuitive	Other	Correct	Other
Baseline	71.5% ($\frac{470}{657}$)	23.4% ($\frac{154}{657}$)	5.0% ($\frac{33}{657}$)	83.3% ($\frac{505}{606}$)	16.7% ($\frac{101}{606}$)
Interruption	65.3% ($\frac{360}{551}$)	29.0% ($\frac{160}{551}$)	5.6% ($\frac{31}{551}$)	79.0% ($\frac{384}{486}$)	21.0% ($\frac{102}{486}$)
Time pressure	60.6% ($\frac{251}{414}$)	27.8% ($\frac{115}{414}$)	11.6% ($\frac{48}{414}$)	67.9% ($\frac{250}{368}$)	32.1% ($\frac{118}{368}$)

Table 1

Percent of trick and no-trick questions answered as correct, intuitive, and other in the baseline, interruption, and time pressure conditions. Parentheses show the number of responses out of total questions per condition.

Such a pattern of confidence ratings, together with the lack of not-intuitive errors, suggests that participants were not aware that their performance on trick questions was affected by the interruption. On the other hand, those under time pressure made more errors, but also understood the impact of that stressor on their ability to solve these problems, reflecting their awareness with confidence ratings.

No-Trick Questions

The responses participants gave to no-trick questions are also consistent with our interpretation thus far. Time pressure had a similar effect on both accuracy and confidence for no-trick questions as it did for trick questions. In particular, we see that, compared to the baseline condition, participants made fewer correct responses when under time pressure (83% and 68%, respectively, $p_{samples} < 0.00001$), and participants seemed to communicate their awareness of this detriment in their confidence ratings ($\gamma = 0.83$ and $\gamma = 0.64$ for baseline and time pressure conditions, respectively, $p_{samples} < 0.00001$).

Unlike time pressure, interruptions had a different effect on behavior during no-trick questions than they did for trick questions. Interrupted participants made fewer correct responses (79%) compared to baseline (83%), and the difference in accuracy may be noteworthy ($p_{samples} = 0.017$). Importantly, participants anticipated an increase in errors for

Figure 2. The top panel plots model estimates of mean proportion correct. Proportion of responses that were correct, intuitive, or other are shaded in the background (note the no-trick questions only has the possibility of correct or other responses). The bottom panel plots model estimates of mean confidence with the average confidence ratings broken down by response type shaded in grey. Error bars in black represent the central 95% of the posterior distribution, while the error bars in grey reflect standard error.

no-trick questions, with interrupted participants giving lower confidence ratings ($\gamma = 0.76$) than those in the baseline ($\gamma = 0.83$) condition, $p_{samples} < 0.00001$. Looking at the grey bar plots in the bottom-right panel of Figure 2, we see that the reduction in confidence was largely driven by lower-confidence incorrect responses. That is, those that were interrupted seemed to know that their performance was impaired for the no-trick questions, unlike the trick questions.

Discussion

We find that both interruptions and time pressure negatively affect people's ability to correctly solve problems. However, the way that time pressure and interruptions affected performance was different, especially for garden-path problems. Participants under time pressure made fewer correct responses for both trick and no-trick questions. For trick questions, they also more often selected the not-intuitive incorrect answers. This not-intuitive answer is not implied by any reasonable understanding of the problem. As such, we expect that the not-intuitive incorrect responses reflect the fact that time pressure sometimes left participants without enough time to read questions carefully, check their answers before making a response, or even establish any clear understanding of the content of the questions. Consistent with this explanation, we saw that those in the time pressure condition gave lower confidence ratings in their responses to both trick and no-trick questions regardless of whether they gave a correct or incorrect response. That is, our results are consistent with the general understanding that time pressure has widespread, deleterious effect on performance, and that people are aware of that fact.

For no-trick questions, the effect of interruptions was similar to time pressure, such that they decreased performance, and participants lowered their confidence as such. Interestingly, interrupted participants seemed to lower their confidence for incorrect responses more than they did for correct responses (though potential scaling issues should lead us to interpret such differences with caution), while those under time pressure lowered their confidence for all responses. Whatever the cause of the errors caused by interruptions for no-trick questions, such confidence ratings suggest that participants were simply not sufficiently motivated to put in the time and effort to recover from the disruption. Outside of trivial experiments like these, such an issue is unlikely to be a problem, at least in cases where people have good reasons to make sure they do not make mistakes (time permitting). Crucially, such affordances may not be possible for trick questions.

Interruptions seemed to have a more problematic effect on trick questions: participants who were interrupted made more intuitive (incorrect) answers, but, unlike the no-trick questions, they did not reduce their confidence. However, we saw that confidence ratings for

intuitive answers were lower than correct responses, suggesting that even in the baseline condition people are not completely oblivious to the errors they have made (De Neys et al., 2013). Importantly, although participants made more errors when interrupted, the pattern of confidence ratings did not change, staying as high as they were in the baseline condition. So, while participants in all conditions seemed to anticipate that they may have made an error, those who were interrupted seemed unaware of the additional impact of interruptions on garden-path problems.

We now discuss one reason why accuracy, but not confidence, for garden-path problems could be affected by interruptions. Recall that an interruption requires the reconstruction of the problem from memory. Such recreation must be based on whatever representation(s) the person had in mind at the time of the interruption. That is, in attempting to bring the problem back into the focus of attention, the person must rely on their memory of what aspects of the problem was important – i.e., their schema for the problem (Bartlett & Bartlett, 1995). In the case of garden-path problems, any understanding of the problem has been intentionally misguided. As such, information that may have been crucial for realising the ‘trick’ may be lost during the reconstruction of the problem. Further, as a result of losing this critical information, the person may also be less likely to realize their schema is problematic and, importantly, may have fewer reasons to doubt their eventual answer. For example, relying on memory to reconstruct the details of a bat-and-ball trick question, we may incorrectly recall that the bat cost \$1, not realising the importance of the clause that the bat cost \$1 *more* than the ball. Should we attempt to verify our answer, the answer that the ball cost 10c would appear correct, thus leading to a high-confidence, intuitive, but ultimately incorrect, answer. Note that if there was no interruption, then such a mistake may still be made. For example, a person may choose not to check their answer against their memory of the question. However, without an interruption it is more likely that the information could be retrieved accurately (e.g., verbatim) from short-term memory stores.

It is reasonable to expect interruptions to cause problems for garden-path situations in the natural world. For example, consider a hypothetical case of a doctor trying to diagnose a patient’s symptoms. An initial examination will reveal a number of facts, their importance

dictated by whatever diagnoses are being considered. In the case where the doctor is fairly confident in a particularly common diagnosis, then only a few details may be deemed relevant. Perhaps the doctor is interrupted while making note of these details. When the primary task is resumed, then the information not considered important at the time of the interruption is less likely to be recalled. Should any of the lost information be relevant for the correct diagnosis, then the performance of the doctor will have been compromised by the interruption (Graber et al., 2014; Johnson et al., 1988; Westbrook et al., 2010). Such a sequence of outcomes seems possible in any garden-path scenario where interruptions may occur, and so our results seem to have far-reaching implications. Interestingly, there may be a relatively easy fix.

Solution. In the absence of time pressure, people have a number of strategies for dealing with the effects of interruptions (e.g., revisiting notes, returning to an earlier stage in a sequence of actions than deemed necessary, asking a colleague for a quick recap of recent events, etc.). We suggest that such strategies may be useful in situations that go beyond what is implied by our current expectations. That is, strategies for recovering from interruptions may be useful even in cases where a person experiences no reason to employ such remedies for an interruption – e.g., when reconstructing the details of a problem appears to happen seamlessly. For example, even if our hypothetical doctor trusts that she was able to recall all important details despite having been interrupted, errors may be avoided by restarting the examination (preemptively taking notes, as verbatim as possible, and using information technology are typically encouraged to reduce the frequency of diagnostic errors (Bates et al., 2001; Graber et al., 2014)). More generally, if people are aware that interruptions can cause them to miss details whose importance may not be immediately obvious, then they will know to take preventative steps whenever interruptions are likely. Hopefully, such action could minimize the problems introduced by interruptions that we identify here.

We expect that errors caused by interruptions could be minimized if people are not only aware that they may make intuitive mistakes, but are also aware that interruptions make such errors more likely. We attempted nudge-like, feedback manipulations to see whether people would realize that there were trick questions without explicit instruction, and react accordingly (see Experiments 2 and 3 in our supplementary materials for details). Our

experiments yielded mixed results, but are consistent with feedback being helpful when it also helped the person become aware of the existence of trick questions. Raelison and De Neys (2019) found clearer evidence of the importance of awareness using a two-response ‘bat-and-ball’ paradigm, where the first response was made under time pressure and the second with no time constraint. Roughly 40% of participants answered the majority of the trick questions correctly on both choices, suggesting that a substantial group of participants understood that trick questions existed. For those who were aware of the trick, whenever they made an incorrect response under time pressure, they were roughly four times more likely to respond correctly on the second attempt than to repeat their error. This suggests that as long as participants are aware of the trick, and understand that there are circumstances that lead them to make mistakes, then they have the ability to fix their mistakes. Though Raelison and De Neys (2019) used time pressure, it seems reasonable to believe that telling people about the potential impact of interruptions on garden-path problems would be a good solution to the issues we outline here, and should be preferred over any nudge-based intervention.

Limitations. Across all three experiments, we used the same, simple addition problems to interrupt people. It is possible that participants in our studies thought the interruption was so simple that it was unlikely to disrupt them from the word problems. This would also explain why participants rarely went back to re-read previous pages of the problem after being interrupted (see supplementary materials for details). A counter-intuitive prediction worthy of future research would be to see if increasing the difficulty of the interrupting task would improve performance on the trick problems. This seems reasonable because if the interrupting task is clearly more disruptive, participants may be more likely to go back and read the problems again and perhaps notice details that they originally missed.

Though we provide one reason for how interruptions impact garden-path problems, the different patterns of accuracy and confidence ratings observed across the questions in our experiment suggest that our explanation is incomplete (see Table 3 in the Appendix). For example, it remains unclear why interruptions seemed to have no impact on some ‘trick’ questions, or why people appropriately reduced their confidence with increased errors for others. In order to have a more fine-grained understanding of the effects of interruptions on

trick problems, we recommend that future research use garden-path problems that are constructed to allow for more targeted interventions such as presenting interruptions at specific moments, or comparing different types of interruptions.

Conclusion

Here, we identified a problem: interruptions made people more susceptible to garden-path problems. We also proposed a potential solution: in situations where we anticipate a garden-path problem, then established strategies for recovering from interruptions should be used even when the person has no immediate reason to think that the interruption was disruptive. We believe this advice is particularly useful in environments where garden-path problems may be more common or may result in serious and potentially devastating consequences.

References

- Altmann, E. M., & Trafton, J. G. (2004). *Task interruption: Resumption lag and the role of cues* (Tech. Rep.). MICHIGAN STATE UNIV EAST LANSING DEPT OF PSYCHOLOGY.
- Bago, B., & De Neys, W. (2019). The smart system 1: Evidence for the intuitive nature of correct responding on the bat-and-ball problem. *Thinking & Reasoning*, 25(3), 257–299.
- Bar-Hillel, M., Noah, T., & Frederick, S. (2018). Learning psychology from riddles: The case of stumpers. *Judgment & Decision Making*, 13(1).
- Baron, R. S. (1986). Distraction-conflict theory: Progress and problems. In *Advances in experimental social psychology* (Vol. 19, pp. 1–40). Elsevier.
- Bartlett, F. C., & Bartlett, F. C. (1995). *Remembering: A study in experimental and social psychology*. Cambridge University Press.
- Bates, D. W., Cohen, M., Leape, L. L., Overhage, J. M., Shabot, M. M., & Sheridan, T. (2001). Reducing the frequency of errors in medicine using information technology. *Journal of the American Medical Informatics Association*, 8(4), 299–308.
- B. Edwards, M., & Gronlund, S. D. (1998). Task interruption and its effects on memory. *Memory*, 6(6), 665–687.
- 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* (pp. 2971–2980).
- Clement, J. (1982). Students' preconceptions in introductory mechanics. *American Journal of physics*, 50(1), 66–71.
- Cohen, H. A. (1975). The art of snaring dragons. artificial intelligence memo number 338. revised.
- De Neys, W. (2020). Rational rationalization and system 2. *The Behavioral and Brain Sciences*, 43, e34–e34.
- De Neys, W., Rossi, S., & Houdé, O. (2013). Bats, balls, and substitution sensitivity:

- Cognitive misers are no happy fools. *Psychonomic Bulletin & Review*, 20(2), 269–273.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic perspectives*, 19(4), 25–42.
- Gillie, T., & Broadbent, D. (1989). What makes interruptions disruptive? a study of length, similarity, and complexity [Journal Article]. *An International Journal of Perception, Attention, Memory and Action*, 50(4), 243–250. doi: 10.1007/BF00309260
- González, V. M., & Mark, G. (2004). "constant, constant, multi-tasking craziness" managing multiple working spheres. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 113–120).
- Graber, M. L., Sorensen, A. V., Biswas, J., Modi, V., Wackett, A., Johnson, S., . . . Singh, H. (2014). Developing checklists to prevent diagnostic error in emergency room settings. *Diagnosis*, 1(3), 223–231.
- Johnson, P. E., Duran, A. S., Hassebrock, F., Moller, J., Prietula, M., Feltovich, P. J., & Swanson, D. B. (1981). Expertise and error in diagnostic reasoning. *Cognitive science*, 5(3), 235–283.
- Johnson, P. E., Moen, J. B., & Thompson, W. B. (1988). Garden path errors in diagnostic reasoning. In *Expert system applications* (pp. 395–427). Springer.
- Johnson, P. E., & Thompson, W. B. (1981). Strolling down the garden path: Error prone tasks in expert problem solving. In *Proceedings of the 7th international joint conference on artificial intelligence-volume 1* (pp. 215–217).
- Keren, G., & Schul, Y. (2009). Two is not always better than one: A critical evaluation of two-system theories. *Perspectives on psychological science*, 4(6), 533–550.
- Kirmeyer, S. L. (1988). Coping with competing demands: interruption and the type a pattern. *Journal of Applied Psychology*, 73(4), 621.
- Mark, G., Gudith, D., & Klocke, U. (2008). The cost of interrupted work: more speed and stress. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 107–110).
- Melnikoff, D. E., & Bargh, J. A. (2018). The mythical number two. *Trends in cognitive sciences*, 22(4), 280–293.

- Pickrell, J. E., McDonald, D., Bernstein, D. M., & Loftus, E. F. (2016). Misinformation effect. *Cognitive illusions: Intriguing phenomena in judgement, thinking and memory*, 406–423.
- Pleskac, T. J., & Busemeyer, J. R. (2010). Two-stage dynamic signal detection: a theory of choice, decision time, and confidence. *Psychological review*, 117(3), 864.
- Primi, C., Morsanyi, K., Chiesi, F., Donati, M. A., & Hamilton, J. (2016). The development and testing of a new version of the cognitive reflection test applying item response theory (irt). *Journal of Behavioral Decision Making*, 29(5), 453–469.
- Puranik, H., Koopman, J., & Vough, H. C. (2020). Pardon the interruption: An integrative review and future research agenda for research on work interruptions. *Journal of Management*, 46(6), 806–842.
- Raoelison, M., & De Neys, W. (2019). Do we de-bias ourselves?: The impact of repeated presentation on the bat-and-ball problem.
- Ratcliff, R., & Rouder, J. N. (2000). A diffusion model account of masking in two-choice letter identification. *Journal of Experimental Psychology: Human perception and performance*, 26(1), 127.
- Ratwani, R. M., Andrews, A. E., McCurry, M., Trafton, J. G., & Peterson, M. S. (2007). Using peripheral processing and spatial memory to facilitate task resumption. In *Proceedings of the human factors and ergonomics society annual meeting* (Vol. 51, pp. 244–248).
- Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of experimental psychology: Learning, Memory, and Cognition*, 21(4), 803.
- Rosen, L. D., Carrier, L. M., & Cheever, N. A. (2013). Facebook and texting made me do it: Media-induced task-switching while studying. *Computers in Human Behavior*, 29(3), 948–958.
- Sloane, J., Donkin, C., Newell, B. R., & Liang, G. (2019). What's lagging in our understanding of interruptions?: Effects of interruption lags in sequential decision-making. In *Cogsci* (pp. 1063–1069).
- Speier, C., Vessey, I., & Valacich, J. S. (2003). The effects of interruptions, task complexity,

and information presentation on computer-supported decision-making performance.

Decision Sciences, 34(4), 771–797.

Thomson, K. S., & Oppenheimer, D. M. (2016). Investigating an alternate form of the cognitive reflection test. *Judgment and Decision making*, 11(1), 99.

Trafton, J. G., Altmann, E. M., Brock, D. P., & Mintz, F. E. (2003). Preparing to resume an interrupted task: Effects of prospective goal encoding and retrospective rehearsal. *International Journal of Human-Computer Studies*, 58(5), 583–603.

Westbrook, J. I., Coiera, E., Dunsmuir, W. T., Brown, B. M., Kelk, N., Paoloni, R., & Tran, C. (2010). The impact of interruptions on clinical task completion. *BMJ Quality & Safety*, qshc–2009.

Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta psychologica*, 41(1), 67–85.

Wickens, C. D. (1991). Processing resources and attention. *Multiple-task performance*, 1991, 3–34.

Appendix

Trick Items

1. Let's pretend you're running a **race**. You pass the person in second place. What place are you now in? (intuitive = first; correct = second; other = third)
2. A farmer had 15 **sheep**. Many of the sheep got very sick. All but 8 died. How many sheep are left? (intuitive = 7; correct = 8; other = 6)
3. Imagine there are three **elves**. Together, they can wrap three toys in one hour. How many elves are needed to wrap six toys in two hours? (intuitive = 6; correct = 3, other = 9)
4. Jerry received the 15th highest **mark** in the class. Jerry also received the 15th lowest mark in the class. How many students are there in the class? (intuitive = 30; correct = 29; other = 31)
5. In an athletics team, tall members are three times more likely to win a medal than short members. This year the team has won 60 medals so far. How many of these have been won by short **athletes**? (intuitive = 20; correct = 15; other = 30)

No-Trick Items

1. A ship had 500 crates of **oranges**. At the first stop, 100 crates were unloaded. At the second stop, 200 more were unloaded. How many crates of oranges were left after the second stop? (correct = 200; other = 100, 300)
2. Sara, Emma, and Sophia embark on a **river** trip. Each of them brings one item: a kayak, a cooler of sandwiches, and a bag of apples. Sara brought the apples and Emma didn't bring anything edible. What did Sophia bring? (correct = cooler of sandwiches; other = bag of apples, kayak)
3. An **expedition** on a mountain climbing trip was traveling with eleven horse packs. Each horse can carry only three packs. How many horses does the expedition need? (correct = 4; other = 3, 5)
4. A **mechanic** shop had five silver cars, two red cars, and one blue car in the garage. During the day, three silver cars and one red car were picked up, and one black car was dropped off. How many silver cars were in the garage at the end of the day? (correct = 2; other = 3, 4)
5. (Attention Check) This question is not like the rest. It is only to check that you are not a robot and paying attention. Please select 3 and then absolutely certain. (correct = 3; other = 1, 2)

Table 2

The stimuli that were used in all three experiments, including trick and no-trick questions. In the paper and the following table, we will refer to each individual item by the word in bold.

	Accuracy (%)			Confidence (0-100)		
	Baseline	Interruption	Timer	Baseline	Interruption	Timer
Expedition	84.9	76.7	63.4	84.4	72.3	58.6
Mechanic	82.1	76.5	75.8	81.0	74.6	56.2
Oranges	94.1	87.8	75.3	91.9	86.6	78.0
River	72.2	72.0	57.1	81.3	80.1	68.6
Athletes	58.9	57.5	47.3	72.8	74.1	57.5
Elves	82.1	73.3	72.5	85.3	78.7	68.5
Mark	53.3	43.8	40.9	81.6	81.2	66.6
Race	92.2	67.3	69.6	98.5	97.7	85.9
Sheep	84.9	82.3	77.4	95.2	94.6	92.1

Table 3

Accuracy and confidence rating for individual questions in each of the three conditions.

	Baseline			Interruption			Time Pressure		
	Correct	Intuitive	Other	Correct	Intuitive	Other	Correct	Intuitive	Other
Expedition	89.6	NA	55.3	83.0	NA	37.1	72.4	NA	34.6
Mechanic	84.8	NA	63.6	82.3	NA	49.6	63.6	NA	33.2
Oranges	93.0	NA	74.2	91.3	NA	52.1	84.9	NA	57.0
River	88.3	NA	63.4	83.5	NA	71.3	72.5	NA	63.3
Athletes	79.5	63.5	61.6	75.8	74.2	52.2	55.3	62.9	46.0
Elves	89.4	68.9	46.3	82.1	73.2	49.8	76.8	51.9	38.5
Mark	91.3	72.8	63.5	89.3	74.8	75.5	81.3	60.5	50.1
Race	98.7	95.5	NA	99.3	94.4	NA	92.5	73.2	42.0
Sheep	96.4	89.1	85.0	97.4	91.9	24.5	97.8	87.1	26.0

Table 4

Confidence ratings broken down by response type (i.e. correct, intuitive, other).

To see the supplementary materials, please visit OSF by clicking [here](#).