Spring 2018

Inferring the Components of Residual Switch Costs

Daniel Thayer
University of Iowa

INFERRING THE COMPONENTS OF RESIDUAL SWITCH COSTS

by

Daniel Thayer

A thesis submitted in partial fulfillment of the requirements
for graduation with Honors in the Psychology

Eliot Hazeltine
Thesis Mentor

Spring 2018

All requirements for graduation with Honors in the
Psychology have been completed.

Toby Mordkoff
Psychology Honors Advisor

INFERRING THE COMPONENTS OF RESIDUAL SWITCH COSTS

by

Daniel Thayer

A thesis submitted in partial fulfillment of the requirements
for graduation with Honors in Psychology

____
Eliot Hazeltine
Thesis Advisor

Spring, 2018

All requirements for graduation with Honors in Psychology have been completed.

____
J. Toby Mordkoff
Psychology Honors Advisor

Abstract

There are many theories that have attempted to explain the underlying mechanisms of task switching. While these theories have furthered the understanding of switch costs, there are inconsistencies that prevent an all-encompassing explanation. Dykstra et al. (in prep) provided evidence which suggests that switch costs may vary depending on the type of task performed. While this finding is interesting, the probability of switches were not equal between tasks. This may have altered switch costs. The present study utilized the inferred switch task from Dykstra et al. (in prep) and adjusted the switch probabilities to match the comparison task. We found that even after corrections were made, the inferred switch cost was still greater than the comparison switch cost. This finding may help clarify current theories of task switching.

Inferring the Components of Residual Switch Costs

Being able to switch between multiple tasks is usually perceived as an effortless endeavor. When studying for an upcoming test, one is easily able to switch from reading a textbook, to checking their phone. However, switching between tasks is not without consequence. Having to go from one thing to another introduces cognitive delays which makes the overall process slower. This delay is known as a switch cost.

Jerslid (1927) was one of the first to look at how individuals switch between tasks. In this experiment, participants were required to complete two different lists. One list had participants continuously adding three to a number. The other list had participants switch between adding and subtracting three. Performance on the alternating list was much slower than on the repetitive list. The delay present when alternating tasks became known as a switch cost. Since Jerslid, many have developed theories that attempt to encapsulate the framework of switch costs. One such theory is task-set inertia.

The theory of task-set inertia proposes that switch costs are a result of proactive interference from the preceding trial (Allport et al., 1994; Allport and Wylie, 2000). Whenever one is switching between tasks, information in relation to the previous task set conflicts with information required for the next trial. This is supported by the fact that when switching from a more dominant task, a smaller switch cost is present in comparison to the opposite scenario. It is theorized that when performing a weaker task, greater inhibition of the dominant task is required. When a switch from the weaker task to the dominant task is engaged, the greater inhibition of the dominant task takes longer to overcome. This interference translates to switch costs. However, not all theories agree with this account. An alternative framework for switch costs is the two stage model of task-set reconfiguration.

In most two stage models of task switching, there is a portion of the switch cost that occurs prior to stimulus onset and another that occurs once a stimulus has been presented. Rogers and Monsell (1995) developed a task switching paradigm that allowed participants to know of a switch before stimulus presentation. They adjusted the interval between the response on each trial and the stimulus onset of the subsequent trial. They found that if the interval continued to increase, the switch cost would be significantly reduced. However, once the response-stimulus interval (RSI) was lengthened to a second, there was no further reduction in switch trial reaction times (RT).

This suggests that a portion of the switch cost reflects time needed to prepare for the next task, as well as an additional process that cannot be prepared prior to stimulus presentation. Specifically, Rogers and Monsell propose that the RSI allows for endogenous reconfiguration of the task set before the next trial. The switch cost that remains reflects an exogenous component of reconfiguration. This refers to stimuli eliciting a habitual response from individuals even if it is in opposition with the intended action on a given trial. This component is known as the residual switch cost. However, the exogenous account is not the only explanation for residual switch costs.

Meiran (2000) developed a model of task switching which suggests that residual switch costs are a result of conflicting responses. When a response is bivalent, this can cause interference between the two possible outcomes. An example of this would be when one key on the keyboard can mean either ‘odd’ or ‘greater than five’. This model suggests that even after an RSI long enough to reconfigure the appropriate response set, all possible responses are still present in working memory. While one response may be more readily available, there is still conflict between the two possible outcomes. Conflict reflects the endogenous process of selecting the appropriate response. This is in opposition to the exogenous theory presented by Rogers and Monsell (1995), demonstrating that it may be unclear as to what comprises the residual switch cost.

Another factor that influences switch costs is the probability of switches. Monsell and Mizon (2006) manipulated the probability of a switch trial in a cued task switching paradigm. They found that as the probability of switch trials increased, the switch cost decreased. Monsell and Mizon propose that participants may have noticed when the probability of a switch was high. If this is the case, it would be more beneficial for participants to always expect a switch. This results in preparation for the alternative task before a switch indicating cue was presented, reducing the switch trial RT. It also leads to an incorrect reconfiguration during repeat trials. This would then require an additional reconfiguration after stimulus presentation, thus increasing repeat RT. These results would lower the overall switch costs. An alternative explanation suggests that participants may have been ‘neutral’ between the two tasks, meaning that neither task was prepared for prior to cue presentation. This would result in preparation being required for both repeat and switch trials. Using this strategy, mappings are updated on every trial, resulting in lower switch costs. Overall, this study demonstrates that the probability of switch trials influences switch costs. It also suggests that there may be multiple strategies for switching between tasks.

While it was originally suggested that switch costs were primarily based on interference from previous trials or task-set reconfiguration (Allport and Wylie, 2000; Rogers and Monsell, 1995), the two theories may not be mutually exclusive. Monsell (2003) proposes that switch costs may be accounted for by both. Reasoning for this stems from the residual switch cost. Monsell states that interference from the previous task set is likely a contributing factor in the residual switch cost. But even that does not provide a complete explanation. The cause of the residual switch cost has been difficult to definitively define. Rogers and Monsell (1995), Meiran (2000), and several others (Allport et al., 1994; Koch and Allport, 2006) all propose different ideas. It could be the case that one of these accounts is correct, but it seems more likely that several of these theories play a role.

Recently, Dykstra et al. (in prep) provided interesting results that may provide insight into the composition of switch costs. In most task switching experiments, participants are required to switch between the dimensions of a bivalent stimulus in which they need to attend. Dykstra et al. (in prep) instead had participants switch between two response mappings. In this study, participants were presented with colored shapes. Only one feature of those colored shapes would be relevant on any given task. Participants were cued at the end of each trial on which response mappings would be needed for the next trial. There were two versions of this response switching task that are of particular interest to the current study. The directed switch task and the inferred switch task.

In the directed switch task, participants were taught two mappings, 'A' and 'B'. Participants were then instructed after each response on which mapping to use for the next trial. This was indicated by an audio cue of either 'A' or 'B'. In the inferred switch task, participants would be given an audio cue that corresponded to the response that was made according to the current mapping that was in use. For example, one mapping had green associated with 'J' and red associated with 'K'. After pressing the 'J' key, if participants heard ‘red’ and there was a red shape present, they would continue using the same mapping. If participants heard ‘green’ after pressing 'J', since the audio cue did not match the stimuli on the screen, this would have been an indicator to switch mappings.

Even though this was not a traditional task switching experiment, Dykstra et al. (in prep) still found switch costs to be present when participants were required to switch between response mappings. A finding of particular interest was that switch costs were greater in the inferred task than in the directed task. While it is tempting to determine that there is something unique about the inferred switch task, there is one difference between the two tasks that may have led to an increase in switch costs. This was the probability of switches.

Dykstra et al. (in prep) had the proportion of switches in the directed task set at 50%. In the inferred task, the proportion of switches was set at 15%. Previous research has found that the lower the probability of a switch in a cued task switching paradigm, the greater the switch cost (Monsell and Mizon, 2006). The fact that the inferred switch task had such a low probability of switches may have been cause for an increase in switch costs.

The aim of the present study is to see if the difference in switch costs are due to the difference in probability of switches. If the inferred switch task still has a greater switch cost, this would imply that something about the task itself may extend switch costs. One thing that is unique about this task is that participants need to determine switch trials based on cue-response mismatch information. This indirect method of switching could influence the underlying components of residual switch costs. It would suggest that there may be several different ways to switch between tasks. If expanded upon, this finding could explain why there are several reliable theories of task switching.

In order to determine this, a couple of changes were made in the present study. One was that the probability of switches was made equal between the two tasks at 33%. This change increases the probability of switches in the inferred task. Results from Monsell and Mizon (2006) would suggest that increasing the probability of switches is likely to lower the average switch cost.

Another change in the current experiment was in regards to Dykstra et al. (in prep) using the same stimulus sets for both the directed and inferred switch tasks. Koch and Allport (2006) suggests that stimuli and response mappings become associated as one performs a task. When response mappings were swapped in the final block, switch costs rose substantially. While these results attest only to similar stimuli being used within an experiment, it seems plausible that this interference may occur between experiments when performed consecutively. In order to avoid the possibility of stimulus based interference, a different stimulus set was used for the directed and inferred tasks. Even with these adjustments made, we expect that the inferred task will have a greater switch cost.

Method

Participants. 70 undergraduate students from the University of Iowa participated in this study. Students were rewarded one hour of class credit upon completion of the experiment. Participants had normal, or corrected to normal vision as well as normal colored vision. Subjects were randomly assigned a condition (Inferred switch first, n = 35; Directed switch first, n = 35).

Apparatus and Stimuli. Stimuli were presented on a 19 inch Dell monitor. Participants were seated at a desk roughly two feet away from the display. RTs and accuracy information were collected using the computer program Matlab based on key presses from a standard keyboard. Headphones were used to present participants with an audio cue after each trial.

The background for all stimuli was black. A random colored shape was presented on each trial in both tasks. The stimuli were placed at the center of the screen. Participants were required to respond to either the number of sides or the color of the presented item. The stimuli used for the color condition were green and red rectangles. The shape condition stimuli were blue triangles and blue pentagons. The type of stimuli used for each task was randomly selected. Whichever stimulus set that was not used in the first task was the stimulus set used for the second task.

Audio tones were used to indicate which stimulus-response mapping to use for the upcoming trial. In the directed switch task, the audio cues were 'A' and 'B'. In the inferred switch task the audio cues were dependent on the stimulus set currently being used. If the color stimulus set was being used, participants would hear either ‘red’ or ‘green’. If the shape stimulus set was in use, participants would hear either ‘three’ or ‘five’.

Design and Procedure. Participants were randomly assigned to one of two groups in this between subjects experimental design. Both groups performed the directed and inferred switch tasks. The only difference between the groups was the order in which they performed the tasks. Participants in condition one would perform the directed switch task, followed by the inferred switch task. Condition two participants would perform the inferred switch task, then the directed switch task. Stimuli sets used for each task were assigned randomly. If the first task used the color stimuli set, then the next task would use the shape stimuli set and vice versa. Both tasks had two mappings that the participants had to learn. Both tasks required participants to randomly alternate between the two response mappings. Before each task began, participants learned the response mappings and the rules for switching between them. 33% of all trials in both tasks were switch trials. In both tasks, stimuli would be presented until a response was made. 350 milliseconds (msec) after the response, a tone would be played for an additional 350 msec. This tone indicated which mapping to use for the next trial. There was then an 800 msec delay until the next stimulus was presented (see Figure 1).

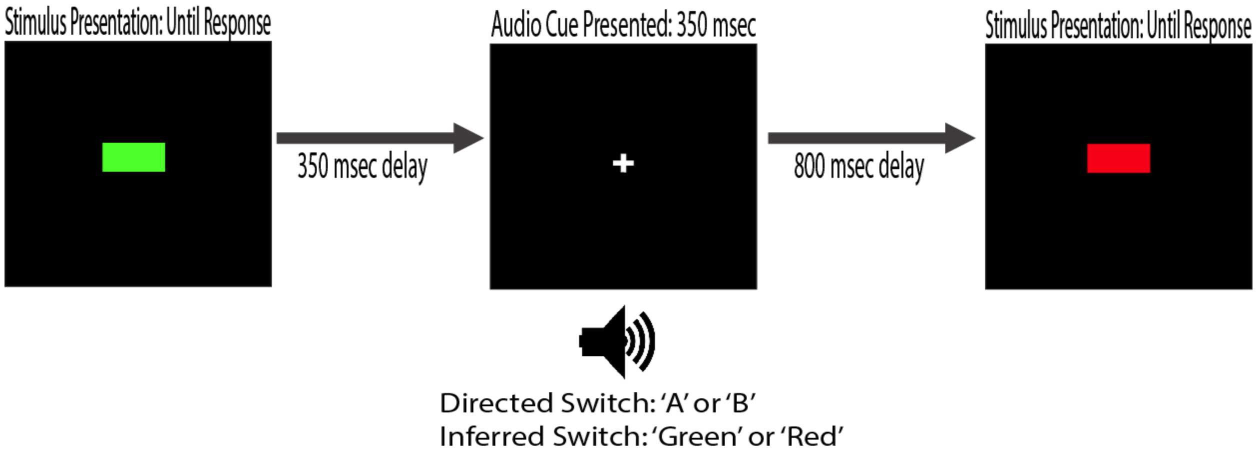

Figure 1. Example of a typical trial when color stimulus set was presented.

The directed switch task had two response mappings. If participants were responding to color, mapping A had the keyboard key 'J' associated with ‘red’ and 'K' associated with ‘green’. When responding to shape, mapping A had 'J' associated with ‘three’ and 'K' associated with ‘five’. Mapping B flipped these pairs. For example, when responding to color, mapping B had 'J' associated with ‘green’ and 'K' associated with ‘red’. Once a response was made, an audible 'A' or 'B' would be played in the headphones. This would inform the participants which mapping to use for the next trial. Participants were instructed to start with mapping A at the beginning of each block. There was a total of 12 blocks in this experiment with 39 trials in each block resulting in a total of a total of 468 trials including practice.

The inferential switch task had a very similar design. In this task, participants were randomly assigned a stimulus type to respond to (shape or color). Participants were taught the same two response mappings described for the directed switch task, but for the alternative stimulus set. How participants would switch between mappings was different in this task. After a response was made, participants would hear a stimulus identifying feature in the headphones. If participants made an accurate response and heard the appropriate identifying feature of the presented stimulus, then the response mapping would stay the same on the next trial. If the audio cue did not match the identifying feature of the presented stimulus, then this indicated a switch in response mappings. For example, on a given trial when the color stimuli set was being used, participants could be presented with a green rectangle. After a response was made, if the audio feedback was ‘green’, this would indicate that the same mapping should be used next trial. If the audio feedback was ‘red’ this meant that that the response mappings had switched. This task had 4 blocks with 60 trials in each block, totally 260 trials including practice.

On 10% of trials, participants would hear a jackhammer noise instead of the usual feedback in the headphones. Participants were told that on trials where this occurred, they were required to maintain the current response mapping. The jackhammer noise was used as an infrequent event in Dykstra et al. (in prep). These trials were maintained in the current study to ensure that it was as consistent to Dykstra et al. (in prep) as possible.

Results

32 participants were excluded from the study because they were not accurate enough on at least one of the two tasks (<.8) or because the program crashed. Additionally, three participants were excluded from condition two to ensure equal group sizes for data analysis. The participants excluded for this reason were randomly selected. This resulted in a total of 38 participants who were used for data analysis (directed switch, n = 19; inferential switch, n = 19). We predicted that the inferential switch task would have a greater switch cost than the directed switch task even when task order, proportion of switches, and stimuli were accounted for.

First, each participant's average RT for switch and repeat trials were calculated. Any trials where the participant gave an incorrect response were removed from data analysis. RT were filtered to only include trials between 200 and 3000 msec. Once average switch and repeat trials were calculated, the difference between the switch and repeat averages were measured. This was the final switch cost for each participant. This method was done separately for both tasks.

In condition one, the average switch cost for the directed switch task was (M = 16.94; SE = 23.01). The average switch cost for the inferred switch task was (M = 376.92; SE = 50.14). In condition two, the average switch cost in the directed switch task was (M = -11.3; SE = 10.19). The average switch cost for the inferred switch task was (M = 465.44; SE = 47.79). A 1(condition) x 2(directed switch: inferred switch) one way analysis of variance (ANOVA) revealed that there was no main effect of condition $F(1,36) = 1.401$, $p = .244$. However, there was a main effect of task $F(1,36) = 119.104$, $p < .001$. This suggests that there is something significantly different between the two tasks. Since there was no significant difference of task switch costs between conditions, we combined the averages of the groups and plotted them in Figure 3.

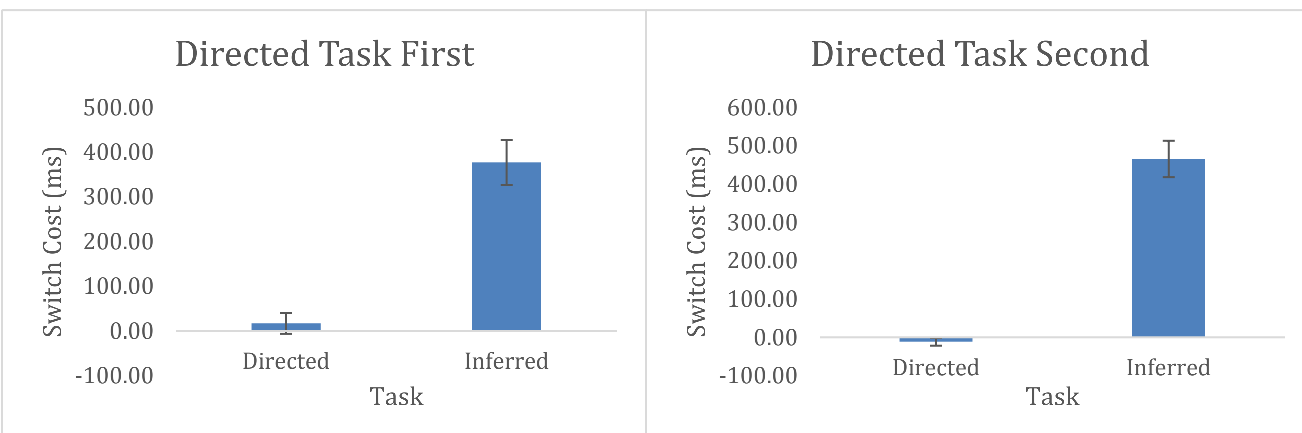

Figure 2. Average switch costs for each task by condition.

Error bars represent standard error of the mean

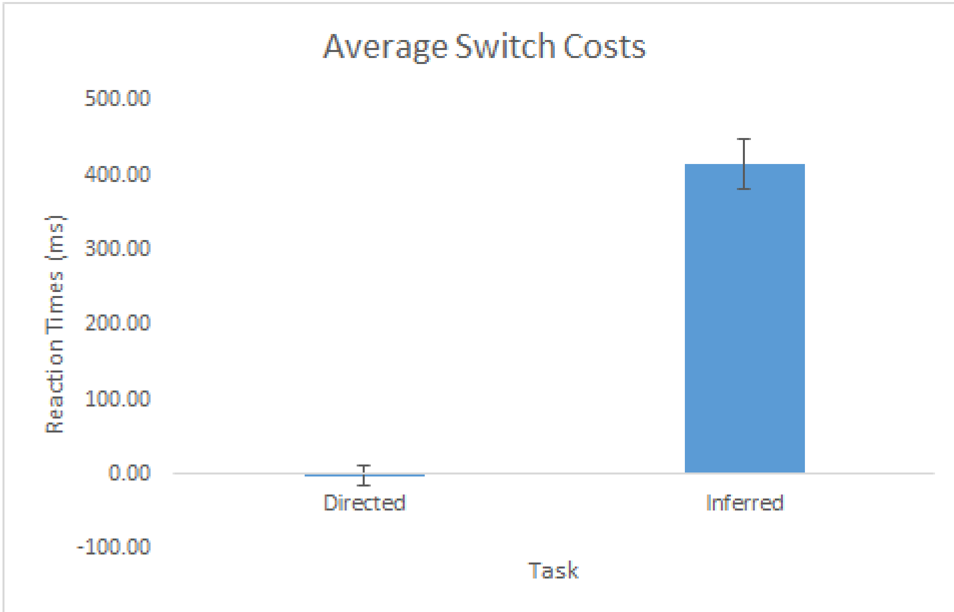

Figure 3. Average switch costs collapsed across conditions.

Error bars represent standard error of the mean

Discussion

The purpose of this study was to determine if the greater switch costs of the inferred task in Dykstra et al. (in prep) reflects something unique about the task itself, or if the results were due to other differences between the tasks. In the present study, we controlled for the probability of switches between the directed and inferred switch tasks. Even after this manipulation, a larger switch cost was still present in the inferred task.

While there was a significant effect of task, it is difficult to determine the precise difference in switch costs between the two tasks. This is due to the fact that there were no switch costs present in the directed task. A speculation as to why this occurred is discussed later. However, the switch costs that were present in the inferred task were relatively large. In fact, the switch costs were greater than those found in Dykstra et al. (in prep). This is particularly interesting because it was expected that the switch costs would be reduced due to the increased probability of switches.

One possible explanation for the increased switch costs of the inferred task in comparison to the directed task is due to working memory. Previous research has indicated that working memory load may have a role in the severity of switch costs (Baddeley et al., 2001). In a typical task switching experiment, participants arguably need to maintain task set information in working memory throughout the experiment. In the inferred switch task, participants are still required to maintain task relevant information, such as response mappings and their relation to presented stimuli. However, switch trials in this task requires the additional process of interpreting cue information. This could constitute additional burden on working memory, therefore increasing the switch costs.

Contrary to this, there is data which suggests that there is minimal interaction between working memory and task switching. Liefooghe (2008) presented evidence that increasing the amount of items that were stored in working memory did not influence switch costs. Although, task switching did interfere with one's ability to maintain information in working memory. While it may be unclear whether working memory load is a component in the inferred task switch costs, it should still be considered a possible explanation.

Response based interference could have also played a role in the increased switch costs. The responses in the inferred switch task were bivalent. Meaning that each possible motor response available to the participants had two potential outcomes. Meiran (2000) states that conflict between response mappings results in interference. This interference is reflected in the residual portion of the switch cost. According to this model of switch costs, the response interference can only be resolved once the stimulus has been presented. The present study had an

RSI of 800 msec. Rogers and Monsell (1995) demonstrated that an RSI greater than 600 msec significantly reduces the endogenous switch cost. If the switch costs found in the current study reflect a similar structure to that proposed by a two stage model, most of it would be reflected by interference in the residual switch cost. Future studies should use univalent responses to help determine if this form of interference is a factor in the overall switch costs of the inferred task.

The results also show that there was no switch cost in the directed task. One interesting aspect of this finding, is that in Dykstra et al. (in prep) the directed task did have switch costs. The only change between the two studies that would have likely altered the switch costs was the probability of switches. It seems unlikely that there were no switch costs in the directed task. Rather, we propose that participants may have been updating their response mappings on each trial.

Due to the nature of this task, the only required information for determining mappings on a subsequent trial is listening to the audio cue. Any other information from the previous trial is irrelevant once a response has been made. Perhaps instead of maintaining response mappings from previous trials, it was more advantageous to update the mappings on every trial. If this is the case, then the additional RT cost of updating the response mapping would be present on both repeat and switch trials. So when subtracting repeat RT’s from switch trial RT’s, no switch costs would be present. This reflects the results of the current experiment. This would indicate that switching task sets does not have to only occur on switch trials. Which would suggest that there are alternative strategies that can be implemented when performing a task switching paradigm.

One possible limitation of this study stems from the large number of participants who were removed from analysis due to low accuracy scores. A majority of the participants who were rejected in this study were excluded due to their low accuracy in the inferred switch task. One possible explanation for this was the difficulty of the task. According to Rubinstein et al., (2001), if a task is difficult, this can increase switch costs. While subjective, we also received feedback from participants suggesting that the inferred task was difficult. Difficulty of task may have been a component of the large switch costs that were observed.

Future research should attempt to further isolate what makes the inferred switch task unique from other switch tasks. If possible, one avenue for future research should investigate if the inferred task switch costs are still present in a paradigm that involves alternating which dimension of a bivalent stimulus to respond to. Further investigation on the details of the inferred switch task may result in a better understanding of the underlying components of switch costs.

References

Allport, A., Styles, E. A., & Hsieh, S. (1994). Shifting intentional set: Exploring the dynamic control of tasks. In C. Umilta & M. Moscovitch (Eds.), Conscious and nonconscious information processing: Attention and performance XV, 421– 452.

Allport, A., & Wylie, G. (2000). Task switching, stimulus-response bindings, and negative priming. *Control of cognitive processes: Attention and performance XVIII*, 35-70.

Baddeley, A., Chincotta, D., & Adlam, A. (2001). Working memory and the control of action: evidence from task switching. *Journal of Experimental Psychology: General*, *130*(4), 641-657.

Dykstra, T.R., Hazeltine, E., & Wessel, J.R. (in prep). *Outcomes of Behavior*. Manuscript in preparation.

Jersild, A. T. (1927). Mental set and shift. Archives of Psychology, 14, 5–81.

Koch, I., & Allport, A. (2006). Cue-based preparation and stimulus-based priming of tasks in task switching. *Memory & cognition*, *34*(2), 433-444.

Liefooghe, B., Barrouillet, P., Vandierendonck, A., & Camos, V. (2008). Working memory costs of task switching. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(3), 478-494.

Meiran, N. (2000). Modeling cognitive control in task-switching. *Psychological research*, *63*(3-4), 234-249.

Monsell, S. (2003). Task switching. *Trends in cognitive sciences*, *7*(3), 134-140.

Monsell, S., & Mizon, G. A. (2006). Can the task-cuing paradigm measure an endogenous task-set reconfiguration process?. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(3), 493-516.

Rogers, R. D., & Monsell, S. (1995). Costs of a predictible switch between simple cognitive tasks. *Journal of experimental psychology: General*, *124*(2), 207-231.

Rubinstein, J. S., Meyer, D. E., & Evans, J. E. (2001). Executive control of cognitive processes in task switching. *Journal of experimental psychology: human perception and performance*, *27*(4), 763-797.