

Buser, Thomas; Peter, Noemi

Working Paper

Multitasking: Productivity Effects and Gender Differences

Tinbergen Institute Discussion Paper, No. 11-044/3

Provided in Cooperation with:

Tinbergen Institute, Amsterdam and Rotterdam

Suggested Citation: Buser, Thomas; Peter, Noemi (2011) : Multitasking: Productivity Effects and Gender Differences, Tinbergen Institute Discussion Paper, No. 11-044/3, Tinbergen Institute, Amsterdam and Rotterdam

This Version is available at:

<https://hdl.handle.net/10419/87018>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

TI 2011-044/3
Tinbergen Institute Discussion Paper



Multitasking: Productivity Effects and Gender Differences

Thomas Buser^{1,2}

Noemi Peter¹

¹ *University of Amsterdam;*

² *Tinbergen Institute.*

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

More TI discussion papers can be downloaded at <http://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 1600

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Duisenberg school of finance is a collaboration of the Dutch financial sector and universities, with the ambition to support innovative research and offer top quality academic education in core areas of finance.

DSF research papers can be downloaded at: <http://www.dsf.nl/>

Duisenberg school of finance
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 8579

Multitasking: productivity effects and gender differences

Thomas Buser Noemi Peter*

February 22, 2011

Abstract

We examine how multitasking affects performance and check whether women are indeed better at multitasking. Subjects in our experiment perform two different tasks according to three treatments: one where they perform the tasks sequentially, one where they are forced to multitask, and one where they can freely organize their work. Subjects who are forced to multitask perform significantly worse than those forced to work sequentially. Surprisingly, subjects who can freely organize their own schedule also perform significantly worse. Finally, our results do not support the stereotype that women are better at multitasking. Women suffer as much as men when forced to multitask and are actually less inclined to multitask when being free to choose.

*Both authors are affiliated with the University of Amsterdam. We are grateful to Robert Dur, Hessel Oosterbeek, Arthur Schram, Joep Sonnemans, Roel van Veldhuizen and seminar participants in Amsterdam for their comments and suggestions. We gratefully acknowledge financial support from the University of Amsterdam through the Speerpunt Behavioural Economics and thank CREED (Center for Research in Experimental Economics and Political Decision-Making) for letting us use their lab. Contact: t.buser@uva.nl, a.n.peter@uva.nl. Website (Thomas Buser): buser.economists.nl.

1 Introduction

There are two common stereotypes about multitasking: juggling tasks is bad for productivity and women are better at it. Surprisingly, both these questions remain scientifically underexplored. This paper fills this gap through an experimental design which allows us to answer the following research questions. First, how does multitasking affect productivity? Second, do people choose the optimal amount of multitasking? Third, are there indeed gender differences in multitasking ability? And fourth, are there gender differences in the propensity to multitask?¹

The first pair of questions is motivated by a practical concern: how to schedule complex tasks optimally? Is sequential execution advisable, or is it more productive to alternate? Is it optimal to let workers choose their own schedule or should companies impose one? Although these questions are important for both workers and companies, economists have traditionally ignored them. The classic paper on multiple tasks (Holmstrom & Milgrom, 1991) is concerned not with scheduling but focuses on asset ownership and incentive contracts between principals and agents.

Fortunately, the importance of work schedules recently gained some attention. Coviello et al. (2010) show that judges who work on many court cases in parallel for presumably exogenous reasons take more time than judges who work sequentially to complete similar portfolios of cases. Their identification strategy uses random variation in the difficulty and urgency of cases (as judged by other judges) and a procedural rule that the first hearing should be held no later than 60 days from filing. They state that “individual productivity cannot be explained only in terms of effort, ability and experience. Individual work scheduling (how much juggling is done) is a crucial input that cannot be omitted from the production function of individual workers” (p. 2).

We share the view that work schedules are an important factor of productivity, but instead of scheduling multiple projects, our research focuses on scheduling multiple tasks. Given that the latter have a much shorter time horizon, we are able to use a lab experiment to examine our research questions. This allows us to randomly allocate subjects to work schedules and thus to circumvent the endogeneity issues that arise when using non-experimental data. Our subjects have to perform two separate tasks according to one of three different treatments: one where they perform the tasks sequentially, one where they are forced to alternate between the two tasks, and one where they can freely organize their work. The amount of time spent on each task is identical in each treatment and performance differences between treatments therefore measure the productivity effect of the different schedules.

There is a related literature on ‘task-switching’ in psychology (see Monsell, 2003 for a review). In these experiments, a series of stimuli is presented to participants who have to perform a short task on each stimulus. For example, pairs of numbers are shown and subjects have to either add them up or to multiply them (see Rubinstein et al., 2001). From time to time, the required operation changes.

¹For the scope of the entire paper, by multitasking we mean switching back and forth between two cognitive tasks. The concepts of multitasking and task-switching are discussed in more detail in Section 2.

It is commonly found that there are ‘switching costs’ associated with changing tasks, i.e. the response to the stimuli is slower after a task-switch. This literature can, however, not answer our research questions adequately. The tasks used are too simple to expect any advantages from multitasking and - related to this problem - subjects are not allowed to choose their schedule freely.² Also, these experiments are not usually incentivized. In contrast, we use two complex tasks of longer duration (a Sudoku and a Word Search puzzle). Therefore subjects can expect an advantage from alternating: they can switch when they get stuck and later look at the same problem with a ‘fresh eye’. Indeed, our subjects do switch when they are allowed to, enabling us to measure the effect of a self-chosen work schedule as well.

Finally, none of the psychological experiments are designed to examine gender differences. Moreover, their samples are generally too small to do so and often characterised by strong gender imbalances. Our comparatively large and balanced sample, on the other hand, allows us to test both whether there are gender differences in multitasking ability and in the propensity to multitask. This pair of research questions is motivated by the gap between popular views and scientific evidence: multi-million selling books advertise that women are better at multitasking as a scientifically established fact³, while in reality this gender difference has not so far been shown by any peer-reviewed experimental paper.⁴ These views are so popular that some biological anthropologists attempt to theoretically explain the empirically-not-yet-established phenomena. Fisher (1999), for example, claims that the prehistoric division of work “built” different aptitudes into the male and female brain through natural selection. Different skills are required for hunting, performed by males, than for gathering, performed by women. As a consequence, argues Fisher, women think “contextually”, as they synthesize many factors into a “web of factors”, while men think linearly, focusing on a single task until it is done. Our design enables us to test whether men indeed follow a more sequential schedule than women when they are allowed to choose freely and whether they perform worse than women when they are forced to multitask.

The paper proceeds as follows: Section 2 clarifies how do we define multitasking, the key concept of the paper. Section 3 explains the details of the experimental design and describes the data. The

²This is true even for papers that claim to examine voluntary task-switching. For example, Arrington & Logan (2004) write that “We examined voluntary task switching by having subjects choose which task to perform on a series of bivalent stimuli. Subjects performed parity or magnitude judgments on single digits. *Instructions were to perform the two tasks equally often and in a random order.*” (p. 610, italics added)

³See for example Pease & Pease (2001) and its adaptation, *Why Men Can Only Do One Thing at a Time and Women Never Stop Talking* (Pease & Pease, 2003).

⁴We searched extensively for peer-reviewed publications about gender differences in multitasking ability, but the closest we could find is Criss (2006) and Havel (2004), which are manuscripts that are made available online at the website of the National Undergraduate Research Clearinghouse. Both examined subjects who had to perform some specified tasks while tallying keywords from a song/story. None of them found gender differences in productivity when multitasking, but Criss (2006) found that women were better at accuracy. Nonetheless, we do not know whether the findings can be attributed to multitasking as none of them had a control group. The media regularly mentions research which supposedly shows that women are better at multitasking but to the best of our knowledge, none of this has been published in peer-reviewed journals.

results are presented in Section 4 while their detailed discussion and the conclusions are presented in Section 5.

2 Definitions

Multitasking is often thought of as the performance of multiple tasks at one time, but this definition is at odds with the findings of many psychologists and neuroscientists. Pashler (1994) reviews the related literature and concludes that our ability to simultaneously carry out even simple cognitive operations is very limited. Using brain scanners, Dux et al. (2006) localize a neural network which acts as a central bottleneck of information processing by precluding the selection of response to two different tasks at the same time. Furthermore, Dux et al. (2009) show that while training can increase the speed of information processing in this brain region, it remains true that the tasks are not processed simultaneously but in rapid succession. Simultaneity is an illusion, which occurs if the tasks are so simple that the alternations are very quick.⁵

Our definition of multitasking is consistent with the above findings: subjects working on two cognitive tasks switch back and forth between them. Our design does not aim for resembling simultaneity but makes the alternation between tasks explicit. Typical situations of this kind are when an employee's work at hand is interrupted with another, perhaps more urgent task. Another example is when people multitask on a computer, switching back and forth between windows or tabs.⁶

Note that our definition is similar to what psychologists call task-switching, but there is an important difference between the two: contingency. In our experiment, subjects *continue* working on the *same problem* after they return from their work on the other task, so they can benefit from a 'fresh eye'. In contrast, in the task-switching experiments subjects get a new stimulus to work on each time (e.g. they get a new pair of numbers to add up, so only the operation remains the same, not the problem they are working on).

3 Experimental design and data

3.1 Treatments and groups

Three treatments were applied during the experiment: Treatment Single, Treatment Multi, and Treatment Choice (subjects were randomly allocated to treatments within each session and they did not know these labels of course). In Treatment Single, subjects had to work on two tasks consecutively,

⁵Simultaneity might be real if at least one of the tasks does not require information processing, but it seems hard to find such a task.

⁶The similarity with the latter is amplified by the fact that the experiment was conducted in a computer lab using a web browser interface.

for 12 minutes each. In Treatment Multi, subjects were forced to switch between the two tasks approximately every four minutes⁷, resulting in the same total time constraint per task as before. Subjects did not know how many switches would occur and the time intervals between switches varied, making anticipation unlikely. In Treatment Choice, subjects could alternate between the two tasks by pressing a ‘Switch’ button, subject to the same time constraint per task as before (12 minutes each). A timer informed subjects about the remaining time for each task. When the 12 minutes for one task expired, the screen changed automatically to the other task and the Switch button could not be used anymore.

It is important to see that this design ensures that the same amount of time is spent on each task in all three treatments. If we tried to resemble simultaneity, for example by splitting the screen, we could not determine how much time subjects spend on each task, and therefore we would not know whether performance between treatments differs due to differences in the amount of time allocated to the two tasks or due to differences in the schedules.

As shown in Table 1, subjects were assigned to three groups. Every subject played two rounds, the first of which was Treatment Single. In the second round, subjects in Group 1 played Treatment Single again, subjects in Group 2 played Treatment Multi, and subjects in Group 3 played Treatment Choice. The subjects knew from the start that there would be two rounds and that they would work on one Sudoku puzzle and one Word Search puzzle in each. The puzzles given in Round 2 were different from the puzzles in Round 1 (but they were the same for all subjects within rounds).

Table 1: Treatments of each group

	Group 1	Group 2	Group 3
Round 1	Single	Single	Single
Round 2	Single	Multi	Choice

This design allows us to answer all four research questions and the fact that Group 1 plays Single twice allows for a difference-in-differences approach. This enables us to correct for learning effects and performance drops due to exhaustion or boredom. To examine the effect of forced multitasking on productivity, we can compare the performance difference between Round 1 and Round 2 of Group 2 to the performance difference of Group 1. To examine the effect of a self-chosen work schedule, we can compare the performance-difference of Group 3 to the performance-differences of the other two groups. If subjects choose the optimal work schedule, we should see that the performance-difference of Group 3 is at least as high as the performance-difference of the other two groups.⁸ Note that subjects already experienced an example of each task in Round 1, so we can assume that subjects in Treatment Choice switch between tasks to maximize their payoff and not due to curiosity.

⁷Gonzalez & Mark (2004) found that information workers spend on average 3 minutes on a task without interruption; this average might be somewhat higher in a less fast-paced environment.

⁸Since subjects in Group 3 can choose whether or not to alternate, finding that they performed worse than the other groups would disprove that they chose optimally.

Figure 1: Sudoku

3					5		6		3	1	4	9	8	5	2	6	7
	5				6	4			8	5	7	3	2	6	4	9	1
6				1				8	6	2	9	7	1	4	3	5	8
		1	8					9	5	3	1	8	4	7	6	2	9
7		2		5	3	8		4	7	9	2	6	5	3	8	1	4
4			2			7			4	8	6	2	9	1	7	3	5
9				3				2	9	6	5	4	3	8	1	7	2
		8	5				4		1	7	8	5	6	2	9	4	3
	4		1					6	2	4	3	1	7	9	5	8	6


To examine gender differences in multitasking ability, we follow a difference-in-difference-in-differences approach. Note that any gender difference in performance can only come from differences in the ability to multitask: since we compare performance in Round 2 to a subject's own performance in Round 1, performance differences cannot be led by differences in task proficiency. Besides, Group 1 captures any gender differences in learning or exhaustion.

Finally, to examine whether there is any gender difference in the propensity to multitask, we check whether there is a gender difference in the number of switches in Treatment Choice. The propensity to multitask might vary with proficiency: subjects who perform well might find switching easier or more beneficial. Alternatively, subjects who get stuck more often may want to switch more often. To avoid attributing such effects to gender differences in multitasking, we control for performance in Round 1.

3.2 Tasks

Our design requires tasks that are not gender-specific and for which multitasking is natural and possibly beneficial. For these reasons, we have chosen Sudoku and Word Search as tasks. Sudoku is played over a 9x9 grid, divided into 3x3 sub-grids called "regions". The left panel of Figure 1 illustrates that a Sudoku puzzle begins with some of the grid cells already filled with numbers. The objective of Sudoku is to fill the other empty cells with integers from 1 to 9, such that each number appears exactly once in each row, exactly once in each column, and exactly once in each region. The numbers given at the beginning ensure that the Sudoku puzzle has a unique solution. For example, the unique solution to the Sudoku in Figure 1 is illustrated the right panel. We measure performance in the Sudoku task by the number of correctly filled cells.

Figure 2: Word Search



When solving a Sudoku puzzle, solutions often come in waves. Multitasking can be appealing when one is stuck: one can work on the other task and hope to see the problem from a different angle when switching back.

The other task was to find as many words as possible in a Word Search puzzle. An example of a Word Search puzzle is presented in the left panel of Figure 2, and its solution is presented in the right panel. Participants had to look for the English names of European and American countries in a 17x17 letter grid. Words could be in all directions, including diagonal and backwards. Subjects' performance is measured by the number of correct words found.⁹

As in the case of Sudoku, it is reasonable to expect subjects to switch when unable to find new words for a while. The situation is similar to polishing a paper, when reading the same lines over and over becomes counterproductive after a while – one changes to another task simply because a 'fresh eye' is needed to recognize meaning behind the letters.

3.3 Procedures, payments, timeline

One pilot and ten regular sessions were run in the computer lab of CREED (Center for Research in Experimental Economics and Political Decision-Making) at the University of Amsterdam. Participants were university students from various fields of study. The application procedure ensured that the two genders were represented approximately equally in every session, but left subjects unaware that the experiment examines gender-related issues. The experiment was conducted in English, therefore both

⁹Subjects did not know in advance how many words were hidden in the puzzle, but they knew that they would be notified once all words were found.

international and Dutch students could participate. All instructions and tasks were computerized,¹⁰ and subjects were not allowed to use any paper or take notes during the experiment.

The experiment started with an introduction that explained the rules of the two tasks and gave the participants opportunity to practice. Subjects learned that there would be two rounds and that they would have to play a Sudoku and a Word Search in both rounds. In each round, subjects earned 6 points for each correctly filled Sudoku cell and lost 6 points for each cell filled with a wrong number to avoid random guessing. Subjects were not penalized for cells filled with multiple numbers.¹¹ They received 9 points for each word found in Word Search. In Word Search, only entire words could be marked and there was therefore no need to penalize random clicking. Subjects' total points for each round were determined as the sum of their points in Sudoku and their points in Word Search. Negative total points were rounded up to 0. One of the two rounds was randomly selected for payment at the end and the conversion rate was 1 euro per 11 points. In addition to this, there was a fixed show-up fee of 7 euros. The performance payments and the conversion rate were chosen based on the results of a pilot, such that subjects could earn approximately equal amounts on the two tasks and that the average payment was around 23 euros. The sessions lasted for approximately 1 hour and 45 minutes.

The order of the tasks within each round was randomized, and the assignment of subjects to the three treatments in round 2 was random as well, so that each group consisted of approximately one third of the subjects in every session. The rules of the treatments were explained immediately before the start of the treatment. Subjects were not aware of the fact that not everyone was playing the same treatment as they did.

After both rounds were over, but before being informed about their payment, we elicited some background information such as gender, age, field of study, and nationality from the subjects through a questionnaire. Those who participated in Treatment Choice were also asked their reasons for (not) switching.

3.4 Data

Our sample consists of 218 subjects from the ten regular sessions.¹² They are 22 years old on average and the majority of them is Dutch (73 percent). Approximately half of the sample consists of economics students (53 percent). The sample contains 11 censored observations from subjects who solved the entire Sudoku puzzle in the second round but not in the first.¹³ As Section 3.1 explained,

¹⁰The program was written in PHP (an HTML-embedded scripting language) and was displayed using the web browser Mozilla Firefox.

¹¹Subjects could enter multiple numbers in one cell to denote uncertainty.

¹²We only use the data from the regular sessions because some parameters were changed after the pilot.

¹³In addition, 17 subjects solved the entire puzzle in the first round and 11 of them also in the second round. These 11 subjects are excluded since we do not know how their performance changed from the first to the second round. We also dropped the six subjects who solved the puzzle only in the first round. Otherwise we would encounter a sample selection problem: among the best performers of Round 1, we would only drop those who fall back enough in Round 2 to not solve the entire puzzle. Inclusion in the sample is thus conditional on not having solved the entire Sudoku in Round 1. Recall

Table 2: Number of observations per cell

	Men	Women	Sample
Group 1	30	40	70
Group 2	39	31	70
Group 3	43	35	78
Total	112	106	218

Table 3: Average total points per cell

	Group 1 (n=70)			Group 2 (n=70)			Group 3 (n=78)		
	Round 1	Round 2	Diff.	Round 1	Round 2	Diff.	Round 1	Round 2	Diff.
Men	184	188	4	186	172	-14	195	174	-22
Women	185	192	7	198	177	-21	205	198	-7
Both	185	190	6	191	174	-17	200	185	-15

Note: all numbers are rounded to the nearest integer.

subjects were randomly assigned to three groups. Table 2 shows the number of observations per group and gender.¹⁴ As we can see, there are between 30 and 43 subjects per cell.

4 Results

4.1 Multitasking and performance

Performance is measured as the sum of Sudoku plus-points and Word Search points.¹⁵ Table 3 shows means per group and gender (for both rounds), and performance differences between rounds. Note that the difference-in-differences(-in-differences) strategy takes care of any performance-differences between cells in Round 1.

Comparing the results of Group 1 and Group 2 to each other shows that the productivity-effects of multitasking are significantly negative: the difference-in-differences is -23 points (ranksum test; p-value=0.02). Subjects who could pick their own schedule (Group 3) perform only slightly better than those forced to multitask and score 21 points less than Group 1 (ranksum test; p-value=0.06).

The difference-in-differences in performance between men and women in Group 2 suggests that men handle forced multitasking relatively better than women, but the difference is not significant (ranksum test; p-value=0.84). The results of Group 3, on the other hand, suggest that women are better at organizing their own schedule, but this difference is not significant either (ranksum test; p-value=0.45). There are no gender differences in learning either: the performance improvement for

that every subject receives treatment Single in Round 1; therefore inclusion is independent of treatment. The proficiency of the dropped subjects is remarkably different from the rest of the sample: they solved all 51 Sudoku-cells in the first round, while the average was 20 cells (with a standard deviation of 11), and the second best performance was 42 cells.

¹⁴The distribution of the dropped subjects is as follows: 5 from Group 1, 8 from Group 2 and 4 from Group 3.

¹⁵Sudoku minus-points were only used to discourage random guessing, not to measure performance.

Group 1 subjects is the same for both genders (ranksum test; p-value=0.87). In sum, simple non-parametric rank-sum tests do not detect any significant gender differences.

Using regression techniques, we can check whether the results hold if we take censoring and the (non-significant) gender differences in learning into account. Table 4 shows the results of fixed effects and first-difference censored regressions which take full advantage of the panel structure of our data.¹⁶ As we can see, the results of the censored regressions are very close to the results of the fixed effect estimates and all the previous conclusions are confirmed. The coefficients of Treatment Multi and Treatment Choice (relative to Treatment Single) are negative and significant at the 5 percent and at the 10 percent level, respectively. The gender-specific estimates confirm that there is no gender difference in learning (the gender dummy is insignificant). The point estimates suggest that men adapted better to Treatment Multi and women adapted better to Treatment Choice, but none of these gender differences is significant.

Table 4: Impact of treatments on total points

	Group-specific estimates		Gender-specific estimates	
	FE	Censored	FE	Censored
Treatment Multi	-22.76** (10.98)	-24.34** (11.44)	-28.39** (12.98)	-31.09** (13.75)
Multi×Male			10.90 (21.78)	13.37 (22.49)
Treatment Choice	-20.97* (11.33)	-21.19* (12.02)	-14.14 (16.30)	-15.51 (17.26)
Choice×Male			-11.63 (23.11)	-9.12 (24.17)
Male			-3.40 (16.39)	-5.49 (17.17)
Nr. of obs.	218	218	218	218

Note: Robust standard errors are shown in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

4.2 Propensity to multitask

To examine gender differences in the propensity to multitask, we use the results of Group 3. Table 5 describes the switching behavior of men and women in Treatment Choice. As we can see, 71 percent of the subjects do actually switch when they are allowed to and the share of switchers is exactly the same for men and for women. So contrary to the claims of Fisher (1999), men do not focus on a single task any more than women do. Moreover, we can reject that women switch more often than men (one-sided t-test; p-value=0.06).

¹⁶Note that since there were two rounds, first difference and fixed effects estimates are equivalent.

Table 5: Number of switches in Treatment Choice

	Men	Women	All
Mean	2.50	1.74	2.16
Standard deviation	2.53	1.67	2.20
Share of switchers	0.71	0.71	0.71
Number of observations	42	35	77

Note: We excluded one subject from this analysis because he misused the ‘Switch’ button (switched multiple times within the same second).

Table 6 displays the results of two OLS regressions where the number of switches is the dependent variable. In Column 1, we only control for performance in Round 1, while in Column 2 we include session and task-order fixed effects. Contrary to our expectations, performance in Round 1 does not influence switching behavior at all; this also implies that the impact of gender on switching is not caused by performance differences. When task order and session fixed effects are also included, the gender difference becomes significant at the 10 percent level. In sum, the results show that if there is any gender difference, it is men switching more than women and not the other way around.

Table 6: Regression results on propensity to switch

Dependent variable: nr. of switches	(1)	(2)
Male	0.76 (0.50)	0.92* (0.53)
Points in Round 1	0.00066 (0.0046)	0.00536 (0.0050)
Nr. of obs.	77	77
Task order and session FE	no	yes

Note: Robust standard errors are shown in parentheses.

Significance levels: *** p<0.01, ** p<0.05, * p<0.1

5 Discussion and conclusions

Our results confirm that work schedules are an important determinant of productivity. We find that multitasking significantly lowers performance in cognitive tasks compared to a sequential execution. This suggests that the costs of switching, which include recalling the rules, details and steps executed thus far, outweigh the benefit of a ‘fresh eye’.¹⁷ Note that this effect differs from the one found by Coviello et al. (2010). In their model, every new task takes resources away from the other active tasks which are closer to being completed, and juggling more tasks consequently increases the *average duration* of task-completion. Our results show that multitasking is bad for productivity even if one is not concerned with average duration.

¹⁷Subjects clearly do expect a benefit from a ‘fresh eye’. Of the majority who chose to switch, many explicitly stated ‘looking at the problem with a fresh eye’ as the main reason.

Subjects who could choose the amount and timing of their switches freely did only marginally better than those forced to switch at unanticipated points in time. Performance under the self-chosen work schedule is actually significantly lower than under the exogenously imposed sequential work schedule. This suggests that subjects fail to choose their own schedule optimally. This finding of inability to organize one's own work optimally is not unprecedented. For example, Ariely & Wertenbroch (2002) find that students who can set their own deadlines perform worse than those forced to adhere to equally spaced deadlines. Another possible explanation is that even though subjects choose the best schedule possible, their performance takes a hit due to the cognitive cost of planning. In a sense, subjects in Treatment Choice had to perform not two but three cognitive tasks: solving the Sudoku, solving the Word Search, and optimizing their schedule. It is difficult to distinguish between these explanations as the number of switches is potentially endogenous to performance.¹⁸ The hypothesis that additional cognitive effort is at the root of the performance impact is however supported by the fact that the average number of switches in Treatment Choice is only 2.16, but subjects still fall back almost as much as subjects in Treatment Multi who were forced to switch four times and could not anticipate the timing of the switches. Whichever explanation is correct, the results are not in favor of self-imposed work schedules.

We do not find any evidence for gender differences in the ability to multitask. Besides, the share of switchers is exactly the same for men and women and the average number of switches is higher for men. Thus, the results contradict the claims of Fisher (1999): if men think so much more linearly than women, why don't they insist more on a sequential schedule? Moreover, why is it that women do not adapt better to multitasking than men when forced to alternate? In sum, the view that women are better at multitasking is not supported by our findings.

References

- Ariely, D. & Wertenbroch, K. (2002). Procrastination, deadlines, and performance: Self-control by precommitment. *Psychological Science*, 13(3), 219–224.
- Arrington, C. M. & Logan, G. D. (2004). The cost of a voluntary task switch. *Psychological Science*, 15(9), 610–615.
- Coviello, D., Ichino, A., & Persico, N. (2010). Too many balls in the air: The impact of task juggling on workers' productivity. *Working paper*.
- Criss, B. R. (2006). Gender differences in multitasking. *National Undergraduate Research Clearinghouse*, 9.

¹⁸Although we find no impact of Round 1 performance, the number of switches might still be endogenous with respect to learning or tiredness effects.

- Dux, P. E., Ivanoff, J., Asplund, C. L., & Marios, R. (2006). Isolation of central bottleneck of information processing with time-resolved fmri. *Neuron*, 52(6), 1109–1120.
- Dux, P. E., Tombu, M. N., Harrison, S., Rogers, B. P., Tong, F., & Marois, R. (2009). Training improves multitasking performance by increasing the speed of information processing in human prefrontal cortex. *Neuron*, 63(1), 127–138.
- Fisher, H. (1999). *The First Sex: The Natural Talents of Women and How They Are Changing the World*. Random House.
- Gonzalez, V. M. & Mark, G. (2004). Constant, constant multi-tasking craziness: Managing multiple working spheres. *CHI Letters*, 6(1), 113–120.
- Havel, M. A. (2004). Gender differences in multitasking abilities. *National Undergraduate Research Clearinghouse*, 7.
- Holmstrom, B. & Milgrom, P. (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics & Organization*, 7, 24–52.
- Monsell, S. (2003). Task switching. *TRENDS in Cognitive Sciences*, 7(3), 134–140.
- Pashler, H. (1994). Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin*, 116(2), 220–244.
- Pease, A. & Pease, B. (2001). *Why Men Don't Listen and Women Can't Read Maps: How We're Different and What to Do About It*. Three Rivers Press.
- Pease, A. & Pease, B. (2003). *Why Men Can Only Do One Thing at a Time Women Never Stop Talking*. Orion.
- 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.