



UvA-DARE (Digital Academic Repository)

Essays in experimental economics

van Veldhuizen, R.R.

Publication date
2013

[Link to publication](#)

Citation for published version (APA):

van Veldhuizen, R. R. (2013). *Essays in experimental economics*. [Thesis, fully internal, Universiteit van Amsterdam].

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 5

Peers at Work in the Lab¹

In an influential study, Mas and Moretti (2009) show that workers are influenced by the productivity of those coworkers who can observe them during the production process. This study attempts to replicate this finding in a laboratory experiment that was specifically designed to capture the fundamental characteristics of the production process of Mas and Moretti (2009). However, we find no evidence of peer effects, which suggests that the findings of Mas and Moretti (2009) may not be as general as often claimed.

5.1 Introduction

Peer effects are potentially relevant in production processes whenever workers interact with coworkers. Introducing a highly productive worker into a shift may for example lead to social spillovers by allowing coworkers to learn better production strategies or encouraging them to increase their effort levels. If such social spillovers play a role, this has important consequences for the choice of production environment (e.g. individual vs team production) as well as the level of worker compensation (i.e. workers should be remunerated based on their production as well as their social spillovers). In recent years, these observations have led to a growing literature that investigates peer effects in production processes both theoretically (starting with Kandel and Lazear (1992)) and empirically (Falk and

¹This chapter is based on Van Veldhuizen, Oosterbeek, and Sonnemans (2012). We are grateful to seminar participants at the University of Amsterdam, the Tinbergen Institute, the 2012 CeDEx-CREED-CBESS Meeting in Norwich and the 2012 ESA World Meeting for helpful comments. Financial support from the University of Amsterdam Research Priority Area in Behavioral Economics is gratefully acknowledged.

Ichino, 2006; Mas and Moretti, 2009).

A particularly influential empirical study is Mas and Moretti (2009, abbreviated as “M&M, 2009” in the remainder of this chapter), which investigates peer effects among cashiers in a large US supermarket chain. The study shows that cashiers increase their effort levels when a highly productive coworker joins their shift. This study is of particular interest, since its data contain the exact spatial orientation of the workers in every shift. This allows the authors to differentiate between two possible peer effect mechanisms: social pressure and prosocial behavior. The former occurs when workers are averse to being caught exerting low effort by a highly productive worker. Prosocial preferences occur when workers are willing to reciprocate a highly productive worker by increasing their own effort. The authors argue that social pressure is only relevant for cashiers who are in the line of sight of the highly productive worker, whereas for prosocial behavior observability is not important. M&M (2009) find that cashiers are influenced by the productivity of a worker that observes them but not by the productivity of workers that they observe, suggesting that peer effects in their sample are driven by social pressure.

By differentiating between two different peer effects mechanisms, M&M (2009) provide a first attempt to open the ‘black box’ of peer effect mechanisms. This is an important first attempt in the literature, which could be relevant for several reasons. Firstly, different peer effect mechanisms may have different policy implications. For example, if peer effects are driven by social pressure, firms should make sure that highly productive workers are able to watch over (and thus put social pressure on) as many of their coworkers as possible. On the other hand, if peer effects are driven by learning effects, firms should try to facilitate knowledge spillovers between workers.

Secondly, different peer effect mechanisms may counteract each other. For example, introducing a productive worker into a shift increases the effort levels of coworkers who are affected by social pressure, but decreases the effort of coworkers who think it is no longer necessary for them to work hard. If both mechanisms are active simultaneously, the net effect is likely to be smaller than the forces of the mechanism separately, which underestimates the potential of peer effects to improve worker productivity if only positive mechanisms are facilitated. Thirdly and relatedly, understanding different peer effect mechanisms makes it possible to more accurately extrapolate empirical results from one sample to another. In

particular, coefficient estimates are more likely to generalize to samples in which similar mechanisms play a role.² For example, M&M (2009)'s coefficient estimates should be more likely to generalize to production settings where social pressure is possible.

However, the degree to which M&M (2009)'s coefficient estimates generalize to other settings also depends on whether the setting of M&M (2009) reflects a larger class of production processes. Many subsequent studies have assumed that this is indeed the case (Ellingsen and Johannesson, 2008; Fehr, Goette, and Zehnder, 2009; Moretti, 2011; Charness and Kuhn, 2011). For example, Ellingsen and Johannesson (2008) state that M&M (2009) shows that “low-productivity workers put in more effort when observed by high-productivity peers.” Thus, M&M (2009) is implicitly assumed to have high external validity.

Whether this assumption is warranted should ultimately be an empirical question, however. Replications of M&M (2009) in other settings could inform us on whether their results can indeed be extrapolated to other contexts. However, in practice such replications are difficult to do because of the level of detail required by the data. Replicating the results of M&M (2009) would require a data set that contains information about the relative spatial positioning of different workers as well as a detailed individual-specific measure of production. When such data are not available in the field, laboratory experiments provide a different means to approximate the key characteristics of M&M (2009). To the extent that the findings of M&M (2009) reflect fundamental aspects of human behavior that also apply outside the original context, they should be replicable in a laboratory environment as well.

In this study, we replicate the setting of M&M (2009) in a laboratory experiment. Our experiment closely follows the set-up of M&M (2009) in a number of important ways. Participants in the role of workers are members of teams and are not financially dependent on the effort of other workers but instead face a higher workload if their team members exert lower effort. Workers perform a real effort task, we vary worker observability and we obtain a direct measure of (permanent) productivity. We vary the extent to which permanent productivity is visible to coworkers in the experiment, which allows us to differentiate between peer effects based on contemporary production and baseline productivity, something that

²See Carrell, Sacerdote, and West (2011) for an interesting case in point.

M&M (2009) are by their own admission unable to do.

In line with M&M (2009), we expected workers to work harder when partnered with more productive coworkers who can observe them during the production process. However, we find that this result does not generalize to the laboratory context. We observe no evidence of peer effects despite the finding that workers are well aware of the production and productivities of their coworkers when this information is available and have the ability to both increase and decrease their effort levels. This suggests that M&M (2009)'s findings may not be as general as has been assumed.

The remainder of this chapter is structured as follows. We first give a brief overview of the results of M&M (2009). We then describe the design and hypotheses of the experiment, which are based on M&M (2009). The next section gives the results, after which we end with a brief discussion.

5.2 Overview of Mas and Moretti (2009)

M&M (2009) estimate peer effects among a two year sample of 394 cashiers working for a US-based supermarket chain. As a measure of productivity, they use the average number of items scanned by a cashier over a 10 minute interval.³ By the authors' admission, this abstracts away from several potentially important aspects of performance (such as quality of service and absenteeism), though it does provide a precise estimate of cashiers' production speeds. For every 10 minute interval, the authors know exactly which workers were on duty and at which cash register they were working. This allows them to identify the workers' spatial orientation, which they use to define the observing and observable set.

At this point, it is useful to first give a few definitions of terms used by M&M (2009). The *focal worker* is the worker in a shift whose behavior is being analyzed. The *observing set* consists of coworkers who are facing the focal worker. The *observable set* consists of coworkers whom the focal worker is facing.

The empirical estimates are based on variations on the following equation

³To be precise, within each 10 minute interval M&M (2009) include only periods when transactions are taking place. Thus, their productivity estimate is the number of items scanned per 10 minutes divided by the time the cashier was involved in a transaction.

(equation 6 in M&M, 2009).

$$\Delta y_{itcs} = \alpha + \beta \Delta \bar{\theta}_{-itcs} + \pi \Delta N_{tcs} + e_{itcs} \quad (5.1)$$

Here, Δy_{itcs} is the production of worker i in 10 minute interval t , at date c , at store s and relative to worker i 's production in the previous 10 minute interval. $\Delta \bar{\theta}_{-itcs}$ is the change in the average permanent productivity of coworkers, ΔN_{tcs} is the change in the number of workers on duty and e_{itcs} is the error term. Peer effects are captured by the coefficient β ; a positive coefficient indicates that workers increase their production speed when the average permanent productivity of their coworkers increases. The key independent variable -average permanent productivity of coworkers- is not directly observed and needs to be constructed using a separate estimation procedure that corrects for possible influences of arbitrary social interactions.

Note that equation (1) assumes that peer effects operate through permanent productivity. However, in practice workers may also be influenced by the contemporaneous effort of their coworkers. M&M (2009) are by their own admission not able to empirically distinguish between these two mechanisms. An attempt to estimate a model that includes contemporaneous effort as an independent variable would have led to a reflection problem (Manski, 1993).⁴ M&M (2009) argue that their estimates are likely a combination of both the contemporaneous and permanent productivity effects.

M&M (2009)'s main findings are presented in the first three columns of table 5.1. The first column shows the results of the primary estimate of equation 1. The coefficient indicates that increasing the average permanent productivity of the coworkers by 10% increases the production of focal worker i by 1.5%. Column 2 shows that these peer effects only appear for changes in the average permanent productivity of coworkers who can observe the focal worker (the observing set). This suggests that peer pressure is more important than prosocial preferences. Column 3 shows that peer effects are specific to workers with below average productivity.

M&M (2009) report several additional estimates. These show inter alia that peer effects persist over time and are larger for workers who occupy registers

⁴The reflection problem is the identification problem that arises when trying to identify a causal relationship between two individuals (or groups) that are observed simultaneously.

Table 5.1: Main Results of Mas and Moretti (2009)

	Dependent Variable:			
	Difference in log productivity of the focal worker between t and $t - 1$			
	(1)	(2)	(3)	(4)
Δ Average coworker permanent productivity	.15 (.02)			
Δ Average coworker permanent productivity in observing set		.17 (.02)		
Δ Average coworker permanent productivity in observable set		.01 (.02)		
Δ Average coworker permanent productivity (below average worker)			.24 (.03)	
Δ Average coworker permanent productivity (above average worker)			.05 (.04)	
Δ Presence worker in observing set				.031 (.003)
Δ Presence worker in observable set				-.030 (.003)
Observations	1,718,052	1,649,916	1,718,052	1,732,941

Notes. This table displays the results of four regressions taken from Mas and Moretti (2009). Columns (1), (2), (3) and (4) are column (1) of table 2, column (1) of table 6, column (1) of table 3 and column (3) of table 6 in Mas and Moretti (2009) respectively. For column (3), the coefficient for above average workers is the total coefficient for above average workers, which is the sum of the two coefficients reported by Mas and Moretti (2009) in column 1 of table 3. The observing set consists of workers who are facing the focal worker. The observable set consists of workers whom the focal worker is facing. The bracketed numbers are standard errors.

in the close proximity of the focal worker. Most relevant for this study are the results reprinted in column 4 of table 5.1, which show that adding a worker to the observing set increases the focal worker's productivity by 3.1%. Interestingly, adding a worker to the observable set decreases the focal worker's productivity by almost the same percentage.

5.3 Experimental Design

The goal of this study is to see if the results of M&M (2009) can be replicated in a laboratory experiment. The experiment consisted of two stages. In the first (or baseline) stage, participant worked alone, allowing us to obtain a measure of their permanent (or baseline) productivity in the absence of peer effects. In the second (or production) stage, participants worked in teams of four, allowing us to investigate the impact of peer effects. In the remainder of this section we first discuss the baseline stage, followed by a discussion of the production stage and a brief comment on the procedure at the end of the experiment.

5.3.1 Baseline Stage

The experiment was computerized using PHP/MySQL. In the experiment, participants had to perform a production task that consisted of adding three two-digit numbers. We chose this task since it is easy to understand, captures the essential features of the production process described by M&M (2009) and results in sizable differences in productivity between participants, which allows us to examine differences between low productivity and high productivity workers. The three numbers appeared on the computer screen together with information about whether the answer to the previous exercise was correct and the cumulative number of successfully completed exercises up to that point. The sequence of numbers used in the exercises was randomly generated before the first session of the experiment, so that it was identical for all participants; we used a separate sequence for the baseline stage and the production stage.

Upon entering the laboratory, participants were welcomed and assigned to a random computer. They received the instructions for the baseline stage of the experiment on screen; the instructions included a single check-up question. After everyone had finished the instructions, the baseline stage started. In the baseline

stage, participants worked individually for 4 minutes and were paid 10 Euro cents for every correct answer they provided.⁵

5.3.2 Production Stage

After the baseline stage, the experiment moved to the production stage, for which participants received additional instructions and check-up questions. After all participants had finished the instructions and check-up questions, the production stage started. In the production stage, there were three different treatments. In two treatments (BASELINETEAM and TEAM), participants were randomly grouped into teams of four. They were told that as a team they had to solve a number of exercises somewhere between 750 and 1150 (the actual number was 829). There was no longer a fixed production time; participants received a fixed fee of 10 euros for their participation in the production stage, regardless of the number of exercises they had solved.

Importantly, they were also told that during the production stage they might receive information about the number of exercises solved by one or more of their teammates. The left part of their computer screen contained an overview of their team. Figure 5.1 gives the team overview used for treatment BASELINETEAM. An arrow going from one participant to another indicates that this participant could see the number of exercises solved by the other in the production stage. For example, in our set-up participant B could see the number of exercises solved by participant A.

The team structure we used allows us to compare four different information perspectives. Participant A knew the number of exercises solved by him could be seen by one team member, whereas participant D knew he could see the number of exercises solved by one team member. Participant B knew he could both see one team member and be seen by another team member and participant C knew he could neither see nor be seen by another participant. The structure of the team remained fixed for the duration of the experiment.

The difference between treatments BASELINETEAM and TEAM is that participants in treatment BASELINETEAM also learned the number of exercises solved by their team members in the *baseline* part of the experiment (their perma-

⁵An English translation of all instructions and two screenshots can be found in the appendix; the original Dutch version of the instructions is available upon request.

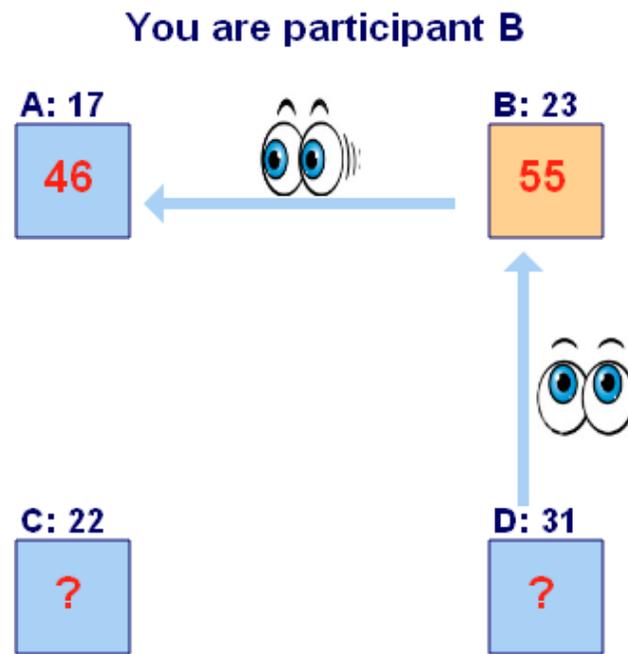


Figure 5.1: Team Overview

Notes. The figure gives the team overview used in treatment BASELINETEAM of the experiment. The numbers above the squares are baseline productivities. These are visible for all team members in treatment BASELINETEAM and are absent in treatment TEAM. The numbers inside the squares are the cumulative production levels in the production stage; these are known only for the participant him/herself and for participants he or she can see (as indicated by the arrows).

nent productivity). They learned the permanent productivity for all participants in their team, even for those for whom they did not know the number of exercises solved in the production stage. As a consequence, treatment BASELINETEAM is closest to M&M (2009) in that it allows peer effects to work through permanent productivity as well as contemporaneous productivity; treatment TEAM only allows peer effects to work through contemporaneous productivity. Note that there is no reflection problem in either treatment, since information flows go only in one direction: for example worker B can be influenced by the contemporaneous productivity of worker A but not vice versa (there is no mirror).

Finally, we also ran an individual treatment in which participants individually had to solve between 188 and 288 exercises (the actual number was 207). Participants in treatment INDIVIDUAL never got any feedback about the performance of other participants in the experiment and were also allowed to leave the experiment after having solved the required number of exercises. We included this treatment to check if organizing workers into teams per se changed their productivity.

5.3.3 End of the Experiment

For each participant, the production stage ended after she (treatment INDIVIDUAL) or she and her team (treatments BASELINETEAM and TEAM) had completed the required number of exercises. After finishing their final exercise, participants received an overview of their earnings and were asked to fill out a questionnaire. The questionnaire contained several demographic questions, a self-monitoring questionnaire (Snyder, 1974) and questions about the experiment. After finishing the questionnaire, participants could collect their payment and leave the laboratory, even if other participants were still solving exercises or working on the questionnaire.

5.4 Hypotheses

M&M (2009) find that increasing the average permanent productivity of the set of coworkers increases workers' production speeds. Importantly, they find that this effect only appears when the change applies to coworkers who are in the focal worker's observing set and that this effect only appears among low-productivity workers. For the experiment, these findings directly translate into the following

hypotheses:

Hypothesis 1 (peer effects): Increasing the average permanent productivity of the set of coworkers increases the production speed of the focal worker.

Hypothesis 2 (observability): Increasing the average permanent productivity of the set of *observing coworkers* increases the production speed of the focal worker; increasing the average permanent productivity of the set of *observable coworkers* has no effect.

Hypothesis 3 (ability): Increasing the average permanent productivity of the set of coworkers increases the production speed of a low productivity focal worker, but not of a high productivity focal worker.

All three hypotheses are testable using treatment BASELINETEAM. Hypotheses 1 and 3 could in principle also apply to treatment TEAM. Although permanent productivity is not directly available to workers in treatment TEAM, workers B and D do know the contemporaneous production of one team member; they could infer the permanent productivity of this team member from his contemporaneous production speed. Hypothesis 2 can only be tested using treatment BASELINETEAM, since permanent productivity levels of observing coworkers are not known and cannot be indirectly inferred in treatment TEAM.

Our design also allows us to investigate the effect of having either one or no coworkers in the observing or observable set. It also makes it possible to examine aspects that lie outside the scope of M&M (2009). Since our set-up avoids the reflection problem, we can check if differences in contemporaneous production of observable workers affect the production speed of the focal worker. We can also investigate to what extent workers are heterogeneous in the way they are influenced by the contemporaneous productivity of their coworkers. By comparing treatments BASELINETEAM and TEAM, we can investigate if knowing the permanent productivity of team members is a necessary ingredient for peer effects to appear. By comparing BASELINETEAM and TEAM with the INDIVIDUAL treatment, we can see if organizing workers into teams per se changes their productivity.

5.5 Results

We ran 9 sessions in February and April 2012, in which a total of 188 subjects participated (84 in TEAM, 84 in BASELINETEAM and 20 in INDIVIDUAL). Participants had an average age of 22.5, 38% reported they studied economics and 58 percent were male.

In the remainder this section, we give an overview of the results. First, we examine evidence of peer effects and investigate if the strength of peer effects is greater for workers in the observing set and for low productivity workers (as per hypotheses 1-3). Second, we investigate several additional effects, some of which could not be estimated using the data of M&M (2009) and look at differences between treatments. Since the results show little evidence of peer effects, we will also investigate possible reasons why peer effects did not appear in our data.

5.5.1 Peer Effects

Hypothesis 1 states that workers' production speed is increasing in the average permanent productivity of their coworkers. As a measure of production speed, we take the average number of exercises solved by the worker per minute in the production stage. As a measure of permanent productivity, we take the number of exercises solved by the worker in the baseline stage.⁶

Table 5.2 shows the result of an OLS regression of worker production speed on average coworker permanent productivity, in which we also correct for the number of exercises solved by the focal worker in the baseline stage. The results in column 1 show that increasing average coworker permanent productivity by 10% increases the production speed of the focal worker by .02%. This is a much smaller percentage than M&M (2009)'s estimate (1.5%, as in column 1 of table 5.1) and is neither an economically nor a statistically significant number. The results in columns 2 and 3 shows that a similar story applies to both treatments taken separately.

⁶As an alternative measure of permanent productivity, we also considered taking only a subset of the baseline, for example only the last two minutes. We chose the overall baseline production since it was most highly correlated ($r=.78$) to production speed in the production stage.

Table 5.2: Peer Effects Estimates

Dependent Variable:			
Log average production speed per minute of the focal worker			
	(1)	(2)	(3)
Log average coworker permanent productivity	.002 (.083)	-.030 (.120)	.011 (.128)
Log focal worker permanent productivity	.678 (.044)	.590 (.065)	.767 (.063)
Constant	.657 (.172)	.841 (.238)	.497 (.288)
Sample	all	BASELINETEAM	TEAM
Observations	168	84	84

Notes. This table displays the results of three OLS regressions; the numbers in parentheses are standard errors.

5.5.2 Observability

Thus we find no evidence that peer effects are relevant at the aggregate level. However, it is possible that this obscures the fact that peer effects are active more locally. Indeed, by hypothesis 2 peer effects should be larger with respect to coworkers who can observe the focal worker. Table 5.3 displays the results of three regressions that examine if this is indeed the case. The regressions examine if the focal worker's production speed is affected by the permanent productivity of the observing coworker, the observable coworker or both. In all cases, we use only the results for BASELINETEAM, since permanent productivity is only available to workers in this treatment.

The results (column 1) show that increasing the permanent productivity of the observing coworker by 10% decreases the production speed of the focal worker by .1%. This effect is much smaller than the effect found by M&M (2009) (1.7%), goes in the opposite direction and is not significant. For observable coworkers, the effect is larger than for M&M (2009) (.48% versus .1%), although it is also not significant and also goes in the opposite direction. Column 3 shows that if both the effect of observing and the effect fo observable coworkers are estimated simultaneously, the resulting estimates are small, negative and not significant as well.

Table 5.3: Peer Effects Estimates by Observability

Dependent Variable: Log average production speed per minute of the focal worker			
	(1)	(2)	(3)
Log permanent productivity (observing set)	-.010 (.122)		-.035 (.261)
Log permanent productivity (observable set)		-.048 (.125)	-.027 (.268)
Log focal worker permanent productivity	.459 (.121)	.508 (.126)	.461 (.247)
Constant	1.04 (.275)	.988 (.283)	1.10 (.674)
Sample Observations	Players A&B 42	Players B&D 42	Player B 21

Notes. This table displays the results of three OLS regressions. The regressions use data from treatment BASELINETEAM only. The numbers in parantheses are standard errors.

5.5.3 Ability

Thus far we have found no evidence of peer effects at the aggregate level or separately for observing or observable coworkers. By hypothesis 3, one reason for the lack of effect could be that peer effects only appear among low productivity workers. To investigate if this is the case, we repeat the regression of table 5.2 separately for low and high productivity workers (using a median split on baseline productivity).

Table 5.4 shows the results of these regressions. Increasing the average permanent productivity of coworkers by 10% increases the production speed of high productivity workers by 1.63%, whereas it reduces the production speed of low productivity workers by .64%. Thus, if anything high productivity workers appear more likely to be affected by peer effects, although neither coefficient is significant at conventional levels. The estimates are similar if we only look at the data from treatment BASELINETEAM (as in columns 3 and 4).

Table 5.4: Peer Effects Estimates by Ability

	Dependent Variable:			
	Log average production speed per minute of the focal worker			
	(1)	(2)	(3)	(4)
Log average coworker permanent productivity	-.064 (.124)	.163 (.103)	-.062 (.164)	.175 (.191)
Log focal worker permanent productivity	.591 (.098)	.846 (.095)	.549 (.130)	.882 (.175)
Constant	2.56 (.929)	-.047 (1.06)	.957 (.337)	-.091 (.538)
Productivity	low	high	low	high
Sample	all	all	BASELINETEAM	BASELINETEAM
Observations	84	84	44	40

Notes. This table displays the results of four OLS regressions; the numbers in parantheses are standard errors.

5.5.4 Additional Effects

Our data also allow us to investigate several additional effects. In particular, M&M (2009) (as reprinted in column 4 of table 5.1) show that adding a worker to the observing set increases production speed, whereas adding a worker to the observable set decreases production speed. In this study, the size of the observable set differs per worker type within each group as per figure 5.1. Using the coefficient estimates of M&M (2009), this suggests that the ranking of production speeds should be $A > B \approx C > D$.

Given the lack of evidence for any peer effects obtained thus far, it is perhaps not surprising that production speeds do not seem to differ systematically between different player types. Figure 5.2 shows that in treatment BASELINETEAM production speeds do not differ systematically between player types. For treatment TEAM, players C produce at a faster speed than players B at the 10% level, but this effect disappears if adjusted for baseline productivity levels.

Using our data we can also investigate the effect of differences in the contemporaneous productivity of coworkers on the production speed of the focal worker. This can be done for the two worker types (B and D) that can see the contemporaneous productivity of at least one coworker. However, table 5.5 shows that the average contemporaneous production speed of the coworker in the

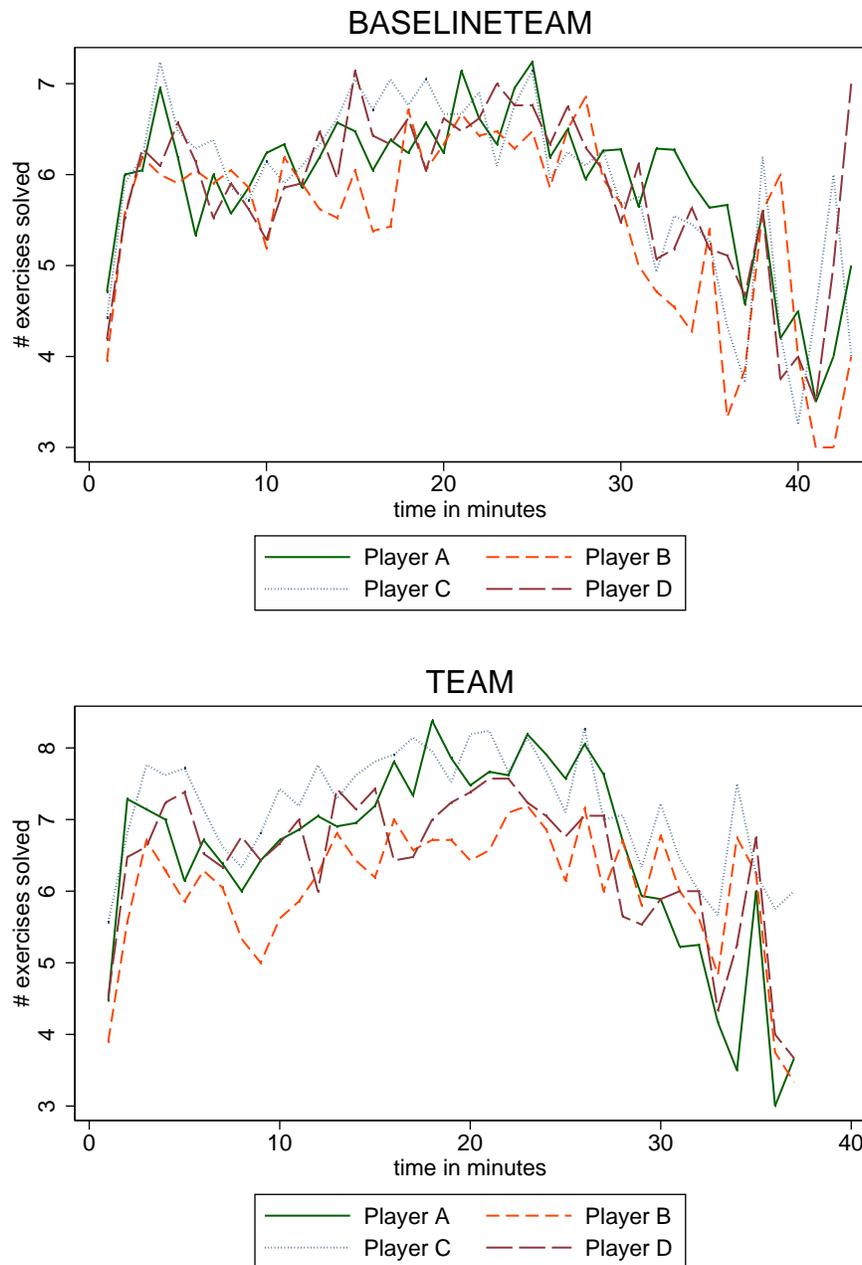


Figure 5.2: Production Speeds

Notes. The figure plots the average number of exercises solved per minute for each player type for treatments BASELINETEAM and TEAM. Note that the declining production speeds after approximately minute 25 are due to a selection effect: slower workers are more likely to remain in the experiment after minute 25.

Table 5.5: Peer Effects Estimates in Contemporaneous Production

Dependent Variable:	
Log average production speed per minute of the focal worker	
Log average contemporaneous production speed (observable set)	-.057 (.080)
Log focal worker permanent productivity	.625 (.082)
Constant	.842 (.194)
Sample Observations	Players B&D 84

Notes. This table displays the results of a single OLS regression. This regression uses data from both treatment TEAM and treatment BASELINETEAM. The numbers in parentheses are standard errors.

observable set does not influence the production speed of the focal worker.

The previous estimate exploits between-worker variation in the production speed of observable coworkers. However, our data also allow us to use within-worker variation in the production speed of coworkers. In particular, the data allow us to see if the focal worker's production speed in minute t is influenced by the production speed of observable coworkers in minute t or minute $t - 1$.⁷ To see if this is the case, we ran a fixed effects regression in which we allowed the focal worker's production speed to depend on the current and lagged production speed of observable workers. Note that these variables are exogenous, since the observable worker does not know the production speed of the focal worker. We also allow for a linear trend to correct for possible learning effects. We also compute the Arellano-Bond estimator to allow the focal worker's production speed in minute t to depend on her production speed in minute $t - 1$.

Table 5.6 shows that the production speed of the focal worker is positively affected by the production speed of observable workers in the current minute and in the preceding minute. The latter effect is significant at the 10% level though not very robust; it disappears if the first minute of data or the coworker production speed minute in t variable is omitted. Thus, there is little evidence that workers are influenced by the contemporaneous production of observable coworkers.

⁷Note that since the number of exercises solved by coworkers was continuously updated, workers were also aware of the production speed of observable coworkers in the same minute.

Table 5.6: Peer Effects Panel Estimates

Dependent Variable: Production speed of the focal worker in minute t		
	(1)	(2)
Production speed in t (observable coworker)	.016 (.019)	.029 (.019)
Production speed in $t - 1$ (observable coworker)	.032 (.019)	.037 (.019)
Production speed in $t - 1$ (focal worker)		-.003 (.024)
Time trend	.019 (.004)	.028 (.006)
Average Fixed Effect	5.67 (.164)	5.42 (.293)
Sample	Players B&D	Players B&D
Observations	2566	2482

Notes. This table displays the results of two regressions. The first column gives the results of a fixed effects regression estimated using OLS, the second column gives the Arellano-Bond estimator. We removed the last minute from the sample since workers did not work for the whole minute. Unlike previous specifications we do not use logs, since several workers did not solve any exercise during one or more minutes. The numbers in parantheses are robust standard errors.

Table 5.7: Peer Effects Estimates by Self-Monitoring

Dependent Variable:				
Log average production speed per minute of the focal worker				
	(1)	(2)	(3)	(4)
Log average coworker permanent productivity	.049 (.254)	-.117 (.203)		
Log permanent productivity (observing set)			-.091 (.193)	.070 (.161)
Log focal worker permanent productivity	.562 (.106)	.712 (.093)	.542 (.173)	.336 (.182)
Constant	.762 (.490)	.836 (.453)	1.04 (.440)	1.10 (.370)
Self-monitoring	low	high	low	high
Sample	all	all	Players A&B	Players A&B
Observations	45	39	24	18

Notes. This table displays the results of four OLS regressions; the numbers in parentheses are standard errors. The regressions use only data from treatment BASELINETEAM

It could also be that the lack of a correlation between the production speed of the focal worker and the production speed of the coworker is due to worker heterogeneity. In particular, a longstanding research tradition in social psychology suggests that people differ in the degree to which they self-monitor to ensure they maintain a desired public image. Highly self-monitoring individuals could be expected to be more susceptible to peer effects, particularly so if social pressure is relevant. However, table 5.7 show that peer effects appear among neither low self-monitoring nor high self-monitoring workers (as measured by the self-monitoring questionnaire of (Snyder, 1974)) even if we only look at the effect of coworkers in the observing set (columns 3-4). Thus, individual heterogeneity in self-monitoring does not seem to explain the lack of peer effects observed in the study.

It could also be that the absence of peer effects on the aggregate is due to other forms of worker heterogeneity. For example, it is possible that some workers decrease their production speed when coworkers work faster, whereas other workers have prosocial preferences and increase their production speed if coworkers produce more quickly. To check this, we re-estimate the regression of column 1 of table 5.6 at the individual level. In 23 out of 84 regressions, either current observable worker production (in 19 cases) or lagged worker production

(in 4 cases) is significant at the 10% level or better, with 19 positive and 4 negative coefficients. The finding that there are only a few negative coefficients suggests that individual heterogeneity is not important. Rather, in line with the results of table 5.6 it appears that only a minority of participants show any evidence of peer effects with respect to observable coworkers.⁸ Thus, we find no evidence that the lack of a significant peer effect in table 5.6 is caused by individual heterogeneity.

5.5.5 Comparing Treatments

So far we have focused on investigating if the findings of M&M (2009) also appear in the laboratory. Comparing treatments TEAM and BASELINETEAM also allows us to see if providing workers with information on the baseline productivity of their peers influences their productivity. Additionally, comparing treatment INDIVIDUAL with both group treatments allows us to see if putting workers in groups rather than making them work individually influences their production speed as well.

Figure 5.3 gives the time series for all three treatments. As column 1 of table 5.8 shows, workers in treatment BASELINETEAM work at a lower production speed than workers in treatment TEAM, whereas the differences with treatment INDIVIDUAL are not significant. However, the difference between TEAM and BASELINETEAM is no longer significant when we correct for differences in permanent productivity (column 2). Column 3 shows that this difference is significant at the 10% level for workers who are not observable even after correcting for baseline productivity, whereas the interaction term in column 4 shows that the treatment difference appears only for workers who are highly productive.⁹ Column 5 shows that only the latter effect remains if both effects are included simultaneously.

5.5.6 Why no Peer Effects?

So far we have seen that peer effects do not seem to play a role in either treatment TEAM or treatment BASELINETEAM. We have also seen that the lack of an

⁸However, note that in table 5.6 the lagged coworker production speed seemed to have a stronger effect, whereas in this case it is coworker production speed in the same minute that is more often significant.

⁹An alternative interpretation of this coefficient is that production speed is less strongly correlated to permanent productivity for workers in treatment BASELINETEAM.

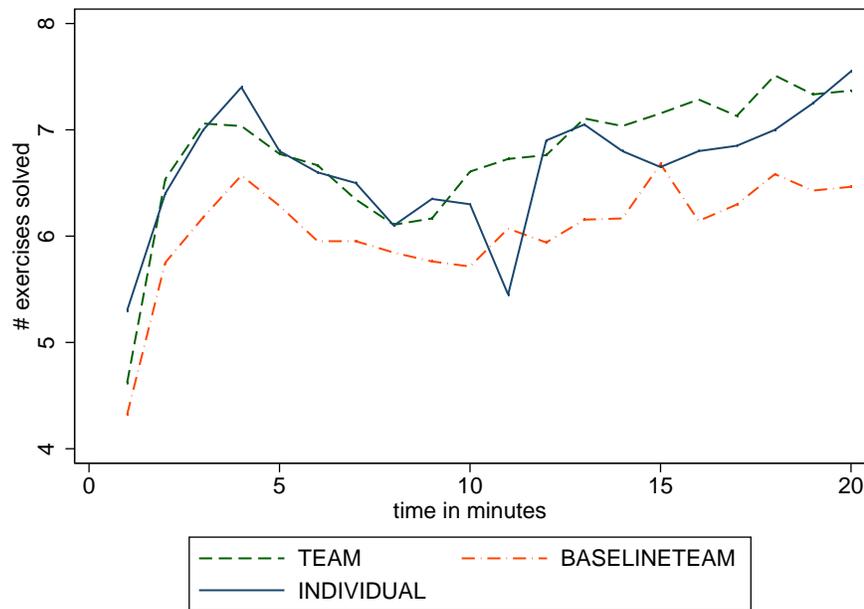


Figure 5.3: Production Speeds by Treatment

Notes. The figure plots the number of exercises solved per minute for treatments BASELINETEAM, TEAM and INDIVIDUAL. The figure displays only the first 20 minutes, since after minute 20 the first participants had finished all exercises, meaning that in subsequent periods the treatments would no longer be directly comparable.

Table 5.8: Peer Effects Estimates by Treatment

Dependent Variable: Log average production speed of the focal worker					
	(1)	(2)	(3)	(4)	(5)
BASELINETEAM	-.114 (.047)	-.036 (.048)	-.076 (.042)	.276 (.147)	.224 (.154)
INDIVIDUAL	-.012 (.075)	-.018 (.030)	-.035 (.053)	-.019 (.048)	-.032 (.052)
Log focal worker permanent productivity		.679 (.042)	.681 (.042)	.769 (.059)	.766 (.059)
Focal worker observable			-.032 (.042)		-.026 (.042)
Focal worker observable x BASELINETEAM			.081 (.060)		.067 (.060)
Log permanent productivity x BASELINETEAM				-.179 (.082)	-.168 (.083)
Constant	1.90 (.033)	.677 (.078)	.689 (.083)	.513 (.108)	.533 (.113)
Sample Observations	All 188	All 188	All 188	All 188	All 188

Notes. This table displays the results of five OLS regressions. The numbers in parantheses are standard errors.

Table 5.9: Recall of Production Stage

Overall production of player				
	A	B	C	D
Player A's estimate	.93	.12	-.26	.07
Player B's estimate	.89	.85	.21	.17
Player C's estimate	-.11	-.26	.93	.27
Player D's estimate	-.12	.70	.44	.91

Notes. This table displays Pearson correlation coefficients. 6 workers did not fill in the recall questions and were thus omitted from the sample, leaving 162 observations.

Table 5.10: Recall of Permanent Productivity

Permanent productivity of player				
	A	B	C	D
Player A's estimate	.94	.82	.48	.71
Player B's estimate	.96	.98	.94	.83
Player C's estimate	.85	.87	1	.82
Player D's estimate	.70	.78	.94	.95

Notes. This table displays Pearson correlation coefficients. 4 workers did not fill in the recall questions and were thus omitted from the sample, leaving 80 observations.

overall effect does not seem to be the result of positive and negative peer effects canceling each other out in the average. In this section, we explore two possible alternative explanations for the lack of peer effects observed in the experiment. Firstly, it could also be that workers are so focused on their own work that they did not pay any attention to the production speeds or permanent productivities of their coworkers. Secondly, it could be that workers are unable to increase or decrease their production speeds.

To investigate the former line of reasoning, we asked workers after the experiment to recall both the overall production and permanent productivities (in BASELINETEAM) of all their coworkers. These recall questions were not incentivized and were not announced until after the experiment had ended. Nevertheless, table 5.9 shows that worker recall of the overall production of observable coworkers was very good (the correlation is between .7 and .9 and significant at the 1% level in all cases). Moreover, table 5.10 shows that workers in BASELINETEAM also did well at recalling the permanent productivity of their coworkers; all correlations are positive, large and significant at the 1% level. This suggests that workers were well aware of both the permanent productivity and (when applicable) the current production of their coworkers. Thus the lack of peer effects observed in

Table 5.11: Variation in Production Speed

Standard deviation of production speed	.451	.495	.425	.372
Weighted standard deviation of production speed	.412	.436	.402	.350
Sample Observations	All 188	BASELINETEAM 84	TEAM 84	INDIVIDUAL 20

Notes. This table gives the average standard deviation of worker production speed relative to the average production speed per exercise. For the second row, the standard deviation is weighted by the production speed of the respective worker relative to the average production speed in the experiment.

the study is not due to a lack of knowledge of the production and productivity of the coworkers.

To see if workers were able to sufficiently adjust their production speed, we investigate if workers were able to change their production speed between exercises. We do not look directly at minutes spent per exercise, since some exercises were more difficult than others. Instead, we look at worker production speed per exercise divided by the average production speed for the respective exercise (among all workers). Table 5.11 shows that the average standard deviation in worker production speed is around 40-50%, which suggests that workers had sufficient scope to change their production speed when required.¹⁰

5.6 Discussion

Overall, the results of Mas and Moretti (2009) do not generalize to the experiment discussed in this study. In particular, we find no evidence of peer effects and also see no evidence that workers are more likely to be influenced by the productivity of coworkers in their observing set. We have also shown that this result is not due to individual heterogeneity and appears despite the finding that workers were well aware of the characteristics of their coworkers and were able to substantially vary their production speed.

It is important to note that we designed the experiment specifically to capture

¹⁰More anecdotally, we observed on several occasions that workers simply stopped working for a few minutes, suggesting that decreasing the production speed was quite possible as well.

the most fundamental elements of the production process discussed by Mas and Moretti (2009). In particular, we used a (not directly incentivized) repetitive real effort task where the number of observable and observing coworkers were carefully controlled. The fact that the experiment fails to replicate their results suggests that the findings of Mas and Moretti (2009) may not be so general as subsequent studies have assumed. In particular, their findings might be specific to certain settings or may be dependent on less fundamental aspects of the production process used in their data (which were not captured by the experiment).

On a more general level, this suggests that there may be an important role for experiments to replicate the results of influential empirical papers. To the extent that empirical results reflect fundamental aspects of human behavior, they should replicate to a wide variety of contexts, including the laboratory. Laboratory experiments can be particularly valuable in cases when it is difficult to replicate a study in the field.

Appendix 5.A Instructions

Welcome to this experiment. During the experiment you are not allowed to communicate with other participants. If you have a question, please raise your hand. One of the experimenters will then come to your cubicle to answer your question.

Today's experiment consists of two parts; part two will take considerably more time than part one. Part two of the experiment will be explained after you have finished part one of the experiment. Your income will be determined on the basis of your results in the experiment. You will also receive a show-up fee of 7 Euros.

Please read through the following instructions carefully. As part of the instructions you will be asked a practice question to test your understanding of the instructions. When you have correctly answered this question, the experiment will move on. Using the navigation bar at the top of your screen it will be possible to return to previous pages during the instructions and practice question.

Instructions Part One

In part one of the experiment the procedure will be as follows. The computer screen will display three two-digit numbers (as in the example below). Your task is to calculate the sum of these three numbers. For every correct answer you will receive 10 Euro cents. An incorrect answer does not earn you any money; after an incorrect answer you will automatically go on to the next exercise. This part of the experiment will take up 4 minutes in total; during these 4 minutes you can do as many exercises as you want. The clock in the lower right corner of the screen tells you how much time you have left. The number of exercises you have answered both correctly and incorrectly is displayed above the current exercise; the (+1) indicates if the previous exercise was answered correctly or incorrectly.

Practice Question

Hank has finished 11 exercises, providing the correct answer to 8 and an incorrect answer to 3. How many Euro cents has Hank earned?

Example of a possible exercise:

Correct Answers: 1 (+1)

Incorrect Answers: 0

What is the sum of the following numbers?

Number A: 16
Number B: 72
Number C: 23

A+B+C=

Your Answer:

Instructions End

You are now ready for part 1 of the experiment. By pressing the link below you will reach a waiting screen. The first part of the experiment starts as soon as all the others have also finished the instructions. On the waiting screen you can read back the text of the instructions.

Instructions Part Two

Like in part one, your task in part two will be to add three two-digit numbers. However, during this part of the experiment you will form a team with three other persons. The experiment will last until you and the other three people in your team have provided a correct answer to a fixed number of exercises. This fixed number of exercises will be somewhere between 750 and 1150 exercises. For this part of the experiment both you and all other team members will get a fixed payment of 10 Euros.¹¹

¹¹These are the instructions for treatment BASELINETEAM. The instructions for the other treatments are available on request.

As soon as you and your team have solved the required number of exercises, the experiment will be over for your team after a short questionnaire. One of the experimenters will come to your cubicle to pay out your earnings for the experiment. After payment you can leave the laboratory, even if the other teams are not done yet.

Information

During the experiment the left side of your screen provides an overview of your team, comparable to the figure displayed below. Each of the squares A, B, C and D represents one of the team members; your square will be colored in orange (in the example below you are participant B). Within the figure, the blue numbers above the squares indicate how many exercises each participant has solved in part 1 of the experiment. Thus, in the example below, participant A has solved 18 exercises, participant B 22, participant C 35 and participant D 21.

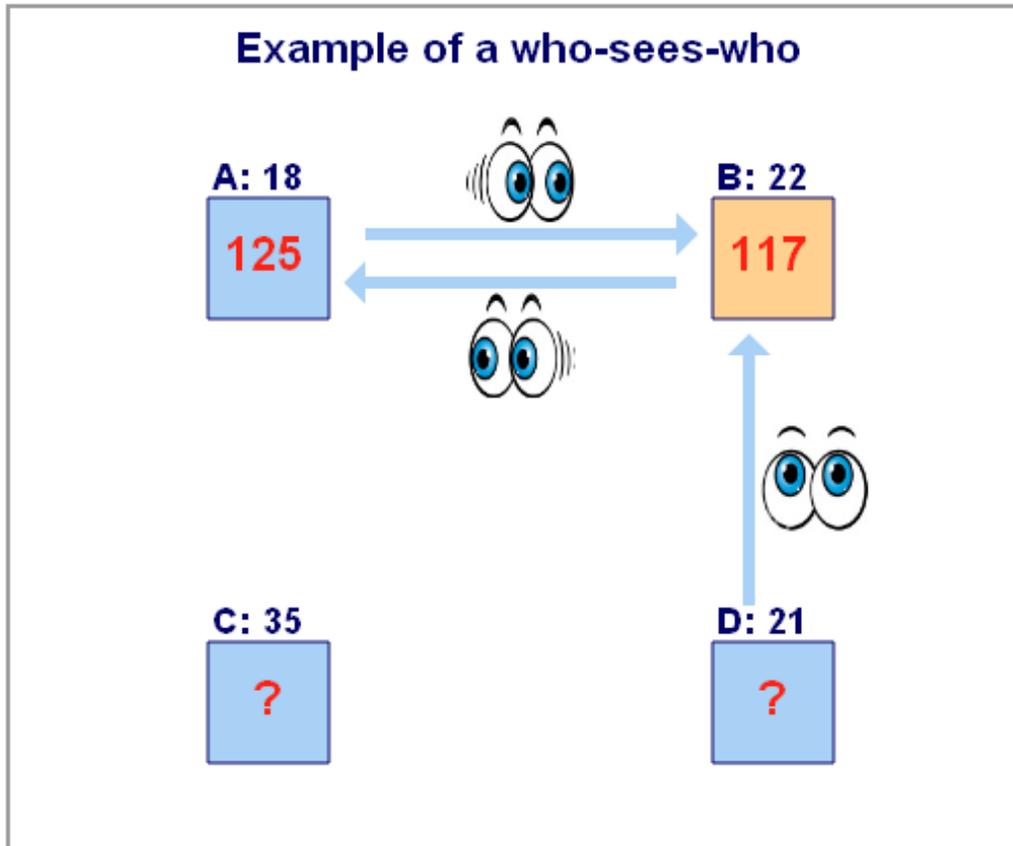
The figure also contains arrows between some participants. In the example below there is an arrow from participant D to participant B, an arrow from participant B to participant A and an arrow from participant A to participant B. These arrows represent information flows. An arrow from one participant to another indicates that this participant is able to see the **number of solved exercises** of the other participant in part two up to that point. Only the number of correct answers will be counted. The number of solved exercises by other participants will be displayed in **red** letters within the square that corresponds to said participant.

In the example below you are participant B and can see how many exercises you have solved (117) and how many exercises participant A has solved (125). In the example below you have **no** information about the number of exercises solved by participants C and D, who therefore have a “?” in their corresponding square. This means that you will at no stage get to know the number of exercises solved by participants C and D (not even after the experiment).

Finally, note that both participant A and participant D can see how many exercises you have solved up to that point. Participant C, however, does not know how many exercises you have solved and will at no stage get to know this number (not even after the experiment).

The figure below only represents an example of a who-sees-who; the who-sees-who that will be used in the experiment (which can have fewer, more or different

arrows) will be announced after the instructions. However, the who-sees-who will remain the same during the experiment; id est, both the arrows, your participant letter and the participant letter of the other participants will remain the same for the whole of the experiment. The composition of your team will not change during the experiment either.



Example Screen

The left part of the screen contains the who-sees-who

On the right part of the screen you do the exercises

If you give a wrong answer, "Oops, incorrect!" will appear on screen during the next exercise

You are participant B

A: 18
125

B: 22
117

C: 35
?

D: 21
?

Oops, incorrect!

What is the sum of the following numbers?

Number A: 23
Number B: 59
Number C: 24

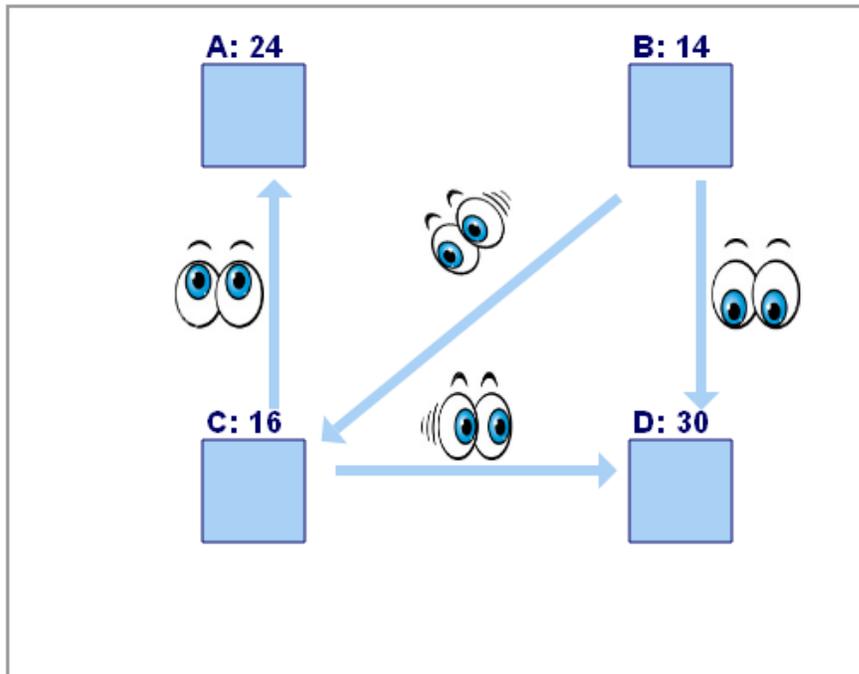
A+B+C=

Your Answer:

Check-Up Question 1

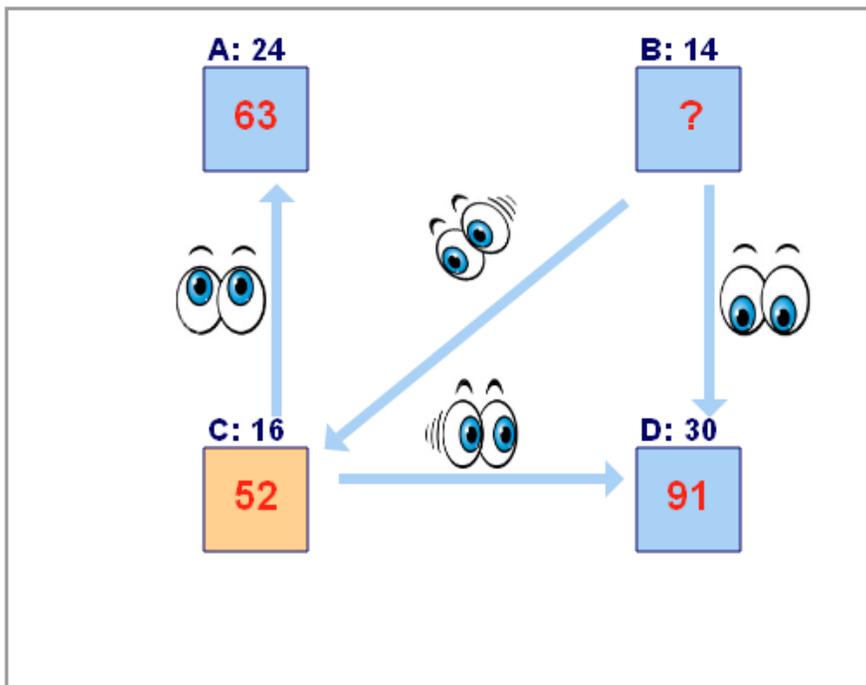
The figure below gives an example of a who-sees-who. Indicate for every team member for which team member he or she can see the number of solved exercises. Also indicate for every team member who can see the number of exercises solved by them.

	Participant A	Participant B	Participant C	Participant D
Participant A knows the number of exercises solved by:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Participant B knows the number of exercises solved by:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Participant C knows the number of exercises solved by:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Participant D knows the number of exercises solved by:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
This participant knows the number of exercises solved by participant A :	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
This participant knows the number of exercises solved by participant B :	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
This participant knows the number of exercises solved by participant C :	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
This participant knows the number of exercises solved by participant D :	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>



Check-Up Question 2

The figure below gives an example of a who-sees-who (the same one as in the previous question). You are participant C, therefore you have all the information that participant C has access to. Indicate for all participants how many exercises they have solved in part one of the experiment. When possible, indicate for every participant how many exercises this participant has solved so far in this part of the experiment (part two). If the number of solved exercises is not known, do not fill in anything.



Check-Up Question 3

Finish the sentence: this part of the experiment ends when *you/your team mates/you and your team mates/everybody in the experiment* have given the correct answer to 750/1150/a fixed number between 750 and 1150 exercises.

Check-Up Question 4

Which participants will know after the experiment how many exercises you have answered correctly?

- All participants in the experiment
- All participants who during the experiment could see the number of exercises you solved.
- All participants of which during the experiment you could see the number of exercises solved.
- Nobody

Instructions End

You are now ready to start part two of the experiment. By pressing the link below you will arrive at a waiting screen. Part two of the experiment will start as soon as all participants have finished the instructions. On the waiting screen you can read back the instructions of this part of the experiment as well.

Appendix 5.B Screenshots

Correct Answers: 1 (+1)
Incorrect Answers: 0

What is the sum of the following numbers?

Number A: 22
Number B: 36
Number C: 75

A+B+C=

Your Answer:

Figure 5.4: Screenshot of the Baseline Phase

You are participant B

A: 17 46

B: 23 55

C: 22 ?

D: 31 ?

Blue: number of solved exercises in part one
Red: number of solved exercises in this part

What is the sum of the following numbers?

Number A: 23
Number B: 59
Number C: 24

A+B+C=

Your Answer:

Submit

Figure 5.5: Screenshot of the Production Phase