



Results from a Replicated Experiment on the Affective Reactions of Novice Developers When Applying Test-Driven Development

Simone Romano^{1(✉)}, Giuseppe Scanniello², Maria Teresa Baldassarre¹,
Davide Fucci³, and Danilo Caivano¹

¹ University of Bari, Bari, Italy

{simone.romano,mariateresa.baldassarre,danilo.caivano}@uniba.it

² University of Basilicata, Potenza, Italy

giuseppe.scanniello@unibas.it

³ Blekinge Institute of Technology, Karlskrona, Sweden

davide.fucci@bth.se

Abstract. Test-Driven Development (TDD) is an incremental approach to software development. Despite it is claimed to improve both quality of software and developers' productivity, the research on the claimed effects of TDD has so far shown inconclusive results. Some researchers have ascribed these inconclusive results to the negative affective states that TDD would provoke. A previous (baseline) experiment has, therefore, studied the affective reactions of (novice) developers—*i.e.*, 29 third-year undergraduates in Computer Science (CS)—when practicing TDD to implement software. To validate the results of the baseline experiment, we conducted a replicated experiment that studies the affective reactions of novice developers when applying TDD to develop software. Developers in the treatment group carried out a development task using TDD, while those in the control group used a non-TDD approach. To measure the affective reactions of developers, we used the Self-Assessment Manikin instrument complemented with a liking dimension. The most important differences between the baseline and replicated experiments are: *(i)* the kind of novice developers involved in the experiments—third-year vs. second-year undergraduates in CS from two different universities; and *(ii)* their number—29 vs. 59. The results of the replicated experiment do not show any difference in the affective reactions of novice developers. Instead, the results of the baseline experiment suggest that developers seem to like TDD less as compared to a non-TDD approach and that developers following TDD seem to like implementing code less than the other developers, while testing code seems to make them less happy.

Keywords: TDD · Affective state · Replication · Experiment

1 Introduction

Test-Driven Development (TDD) is an incremental approach to software development in which unit tests are written before production code [1]. In particular, TDD promotes short cycles composed of three phases to incrementally implement the functionality of a software:

Red Phase. Write a unit test for a small chunk of functionalities not yet implemented and watch the test fail;

Green Phase. Implement that chunk of functionalities as quickly as possible and watch all unit tests pass;

Refactor Phase. Refactor the code and watch all unit tests pass.

Advocates of TDD claim that this development approach allows improving the (internal and external) quality of software as well as developers' productivity [8]. However, research on the claimed effects of TDD, gathered in secondary studies, has so far shown inconclusive results (*e.g.*, [15]). Such inconclusive results might relate to the negative affective states that developers would experience when practicing TDD (*e.g.*, [8]). For example, frustration due to spending a large amount of time in writing unit tests that fail, rather than immediately focusing on the implementation of functionality. Nevertheless, only Romano *et al.* [21] has studied through a controlled experiment the affective reactions of developers when applying TDD to implement software. In particular, they recruited 29 novice developers who were asked to carry out a development task by using either TDD or a non-TDD approach. At the end of the development task, the researchers gathered the affective reactions to the development approach, as well as to implementing and testing code. To this end, Romano *et al.* used Self-Assessment Manikin (SAM) [3]—a lightweight, but powerful self-assessment instrument for measuring affective reactions to a stimulus in terms of the pleasure, arousal, and dominance dimensions—complemented with the liking dimension [17]. The results highlight differences in the affective reactions of novice developers to the development approach, as well as to implementing and testing code. In particular, novice developers seem to like TDD less as compared to a non-TDD approach. Moreover, novice developers following TDD seem to like implementing code less than those developers following a non-TDD approach, while testing code seems to make TDD developers less happy.

The Software Engineering (SE) community has shown a growing interest in replications of empirical studies (*e.g.*, replicated experiments) and recognized the key role that replications play in the construction of knowledge [25]. To validate the results of the experiment by Romano *et al.* [21] (also called baseline experiment from here on), we conducted a replicated experiment with 59 novice developers. In the replication, we investigated the same constructs as the baseline experiment, but in a different site and with participants sampled from a different population—*i.e.*, 59 second-year vs. 29 third-year undergraduates in Computer Science (CS) from two different universities.

Paper Structure. In Sect. 2, we report background information and related work. The baseline experiment is summarized in Sect. 3. The replication is outlined in Sect. 4. The results of our replication are presented and discussed in Sect. 5 and Sect. 6, respectively. We discuss the threats to validity of our replication in Sect. 7. Final remarks conclude the paper.

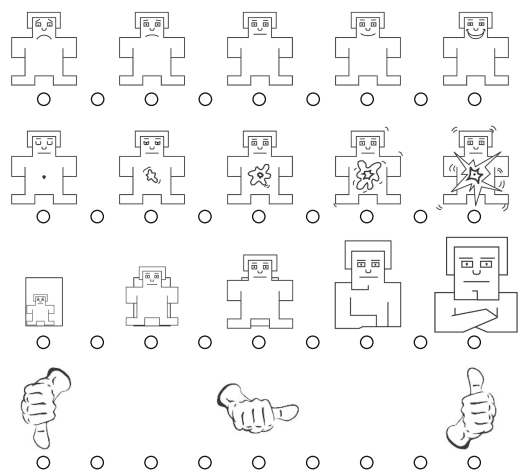


Fig. 1. From top down, the graphical representations of the pleasure, arousal, dominance, and liking dimensions. This figure has been taken from [21].

2 Background and Related Work

According to the PAD (Pleasure-Arousal-Dominance) model—a psychological model to describe and measure affective states—, people’s affective states can be characterized through three dimensions: pleasure, arousal, and dominance [22]. The pleasure dimension varies from unpleasant (*e.g.*, unhappy/sad) to pleasant (*e.g.*, happy/joyful), the arousal one ranges from inactive (*e.g.*, bored/calm) to active (*e.g.*, excited/stimulated), and finally, the dominance dimension varies from “without control” to “in control of everything” [17]. To measure a person’s affective reaction to a stimulus in terms of the pleasure, arousal, and dominance dimensions, Bradley and Lang [3] proposed a pictorial self-assessment instrument they named SAM. This instrument represents each dimension graphically with a rating scale placed just below the graphical representation of each dimension so that a person can self-assess her affective reaction in terms of that dimension (see Fig. 1). For instance, SAM pictures the pleasure dimension through manikins varying from an unhappy manikin to a happy one; thus the nine-point rating scale, placed just below the graphical representation of the pleasure dimension, allows a person to self-assess, from one to nine, that dimension of her affective reaction. Recently, Koelstra *et al.* [17] have complemented SAM with

the liking dimension ranging from dislike—pictured through a thumb down—to like—pictured through a thumb up (see Fig. 1).

Both Human-Computer Interaction (HCI) and affective computing research fields have utilized SAM in their empirical studies (*e.g.*, [12, 17]). Later, the SE research field has used SAM as well. For example, Graziotin *et al.* [11] conducted an observational study with eight developers who performed development tasks on individual projects. Every ten minutes, the participants self-assessed both their affective state, by using SAM, and their productivity. The results show that pleasure and dominance are positively correlated with productivity.

A few SE studies have investigated the affective states of developers through controlled experiments (*e.g.*, [16, 26]). Besides the study by Romano *et al.* [21], which we summarize in the next section, no controlled experiment has been conducted to investigate the affective reactions of developers while practicing TDD.

3 Baseline Experiment

In this section, we summarize the baseline experiment by Romano *et al.* [21] by taking into account the guidelines for reporting replications in SE [6].

3.1 Research Questions

The baseline experiment aimed to answer the following Research Question (RQ):

RQ1. Is there a difference in the affective reactions of novice developers to a development approach (*i.e.*, TDD vs. a non-TDD approach)?

The aim of RQ1 was to understand the affective reactions that TDD raises on novice developers in terms of pleasure, arousal, dominance, and liking. To deepen such an investigation, two further RQs were formulated and studied:

RQ2. Is there a difference in the affective reactions of novice developers to the implementation phase when comparing TDD to a non-TDD approach?

RQ3. Is there a difference in the affective reactions of novice developers to the testing phase when comparing TDD to a non-TDD approach?

The aim of RQ2 and RQ3 was to understand the effect of TDD on the affective reactions of novice developers—in terms of the pleasure, arousal, dominance, and liking dimensions—with respect implementing and testing code, respectively.

3.2 Participants and Artifacts

The participants in the baseline experiment were 29 third-year undergraduates in CS at the University of Basilicata (Italy). According to previous work (*e.g.*, [13]), Romano *et al.* considered undergraduates in CS as a proxy of novice developers. The participants were taking the SE course when they voluntarily accepted to take part in the experiment. Once the students accepted to participate, they were

asked to fill in a pre-questionnaire (*e.g.*, to collect information on their experience on unit testing). Based on the data gathered through this questionnaire, the participants had experience in both C and Java programming. No participant had experience with TDD at the beginning of the SE course.

The baseline experiment used two experimental objects—*i.e.*, Bowling Score Keeper (BSK) and Mars Rover API (MRA). Each participant dealt with either BSK or MRA. The participants, who received BSK, were asked to develop an API for calculating the score of a bowling game, while those who received MRA had to develop an API for moving a rover on a planet. In both cases, they had to code in Java and write unit tests by using JUnit. At the beginning of the experimental session, any participant was provided with: (i) a problem statement regarding the assigned experimental object; (ii) the user stories to be implemented (*i.e.*, 13 user stories for BSK and 11 user stories for MRA); (iii) a template project for the Eclipse IDE containing the expected API and an example JUnit test class; and (iv) for each user story an acceptance test suite to simulate customers' acceptance of that story. Both BSK and MRA had been previously used as experimental objects in empirical studies on TDD and could be fulfilled in a three-hour experimental session (*e.g.*, [9, 10]).

To gather the affective reactions of the participants, Romano *et al.* exploited SAM [3] complemented with the liking dimension [17]. SAM allows measuring people's affective reactions to a stimulus over nine-point rating scales in terms of pleasure, arousal, dominance, and liking (see Sect. 2).

3.3 Variables and Hypotheses

The baseline experiment compared the affective reactions of two different groups of novice developers, namely *treatment* and *control*. The treatment group consisted of participants who were asked to use TDD to carry out a development task, while the control group consisted of participants who were unaware of TDD and had to perform a development task by using a non-TDD approach named YW (Your Way development)—*i.e.*, the approach they would normally utilize to develop [9]. Therefore, the main Independent Variable (IV), or main factor, manipulated in the baseline experiment was **Approach**, which assumed two values: TDD or YW. Within each group, some participants dealt with BSK, while others dealt with MRA. Thus, there was a second IV, namely **Object**, which had BSK or MRA as the value.

To measure the pleasure, arousal, dominance, and liking dimensions with respect to the development approach (*i.e.*, to answer RQ1), Romano *et al.* used the following four ordinal Dependent Variables (DVs): APP_{PLS}, APP_{ARS}, APP_{DOM}, and APP_{LIK}. These variables assumed integer values in between one and nine since each dimension could be assessed through a nine-point rating scale (see Sect. 2). Similarly, they measured pleasure, arousal, dominance, and liking with respect to the implementation and testing phases (*i.e.*, to answer RQ2 and RQ3) through the following four ordinal DVs each: IMP_{PLS}, IMP_{ARS}, IMP_{DOM}, IMP_{LIK}, TES_{PLS}, TES_{ARS}, TES_{DOM}, and TES_{LIK}.

To answer the RQs, the following parameterized null hypothesis was tested:

H0_{DV}. There is no effect of Approach on $DV \in \{\text{APP}_{\text{PLS}}, \text{APP}_{\text{ARS}}, \text{APP}_{\text{DOM}}, \text{APP}_{\text{LIK}}, \text{IMP}_{\text{PLS}}, \text{IMP}_{\text{ARS}}, \text{IMP}_{\text{DOM}}, \text{IMP}_{\text{LIK}}, \text{TES}_{\text{PLS}}, \text{TES}_{\text{ARS}}, \text{TES}_{\text{DOM}}, \text{TES}_{\text{LIK}}\}$.

3.4 Design and Execution

The design of the baseline experiment was $2 * 2$ factorial [27]. Such a kind of between-subjects design has two factors (*i.e.*, two IVs) having two levels each. The two factors were Approach and Object. Each participant in the baseline experiment was randomly assigned to one development approach and to one experimental object—*i.e.*, no participant used both development approaches or dealt with both experimental objects. In particular, 15 participants were assigned to TDD—7 with BSK and 8 with MRA—, while 14 participants were assigned to YW—7 with BSK and 7 with MRA.

Before the experiment took place, the participants had undergone a training period. In the first part of the training period, all participants attended face-to-face lessons on unit testing, JUnit, Test-Last development (TL), and Incremental Test-Last development (ITL). They also practiced unit testing with JUnit in a laboratory session. In the second part of the training, the participants in the treatment group learned TDD and practiced it through two laboratory sessions and three homework assignments. The participants in the control group did not learn TDD, rather they practiced TL and ITL through two laboratory sessions and three homework assignments. Regardless of the experimental group, the assignments were the same. The researcher conducted the experiment in a single three-hour laboratory session at the University of Basilicata where, based on the experimental groups, the participants carried out the development task—*i.e.*, they tackled MRA or BSK—by using TDD or YW. At the end of the development task, the participants were asked to self-assess their affective reactions to the used development approach through SAM [3] complemented with the liking dimension [17]. Similarly, they self-assessed their affective reactions to implementing and testing code, respectively.

3.5 Data Analysis and Results

Romano *et al.* analyzed the effects of Approach, Object, and their interaction (*i.e.*, Approach:Object) by using ANOVA Type Statistic (ATS) [4], a non-parametric version of ANOVA recommended in the HCI research field to analyze rating-scale data in factorial designs [14] (like the case of the baseline experiment). In particular, for each DV, the following ATS model was built: $DV \sim \text{Approach} + \text{Object} + \text{Approach} : \text{Object}$. To judge whether an effect was statistically significant, the α value was fixed (as customary) at 0.05. That is, an effect was deemed significant if the corresponding p-value was less than α . To quantify the magnitude of the effect of Approach, in case it was significant, Romano *et al.* used Cliff's δ effect size [7]. The size of an effect is deemed: *negligible*, if $|\delta| < 0.147$; *small*, if $0.147 \leq |\delta| < 0.33$; *medium*, if $0.33 \leq |\delta| < 0.474$; or *large*, otherwise [20].

In Table 1, we report the ATS results of the baseline experiment. These results show a significant effect of Approach on APP_{LIK} (p-value = 0.0024), namely there is a significant difference between TDD and YW with respect to APP_{LIK} . This allowed rejecting $H0_{APP_{LIK}}$. The difference in the APP_{LIK} values was in favor of YW and large ($\delta = 0.6048$).¹ Accordingly, Romano *et al.* concluded that developers using TDD seem to like their development approach less than those using a non-TDD approach (*i.e.*, answer to RQ1). Table 1 also shows two further significant effects, one for IMP_{LIK} (p-value = 0.0396) and one for TES_{PLS} (p-value = 0.0178) so allowing rejecting $H0_{IMP_{LIK}}$ and $H0_{TES_{PLS}}$, respectively. Both effects were in favor of YW. The effect size was medium ($\delta = 0.4286$) for IMP_{LIK} , while large for TES_{PLS} ($\delta = 0.5$). Based on these results, Romano *et al.* concluded that: developers using TDD seem to like the implementation phase less than those using a non-TDD approach (*i.e.*, answer to RQ2); and the testing phase seems to make developers using TDD less happy as compared to those using a non-TDD approach (*i.e.*, answer to RQ3). As for the effects of Object and Approach:Object, they were in no case significant—*i.e.*, neither the experimental object nor the interaction with the development approach seems to influence the affective reactions of novice developers.

Table 1. Results, from statistical inference, of the baseline experiment.

DV	IV			Cliff's δ	Outcome for $H0_{DV}$
	Approach	Object	Approach:Object		
APP_{PLS}	0.1615	0.7721	0.8998	-	$H0_{APP_{PLS}}$ not rejected
APP_{ARS}	0.2774	0.7794	0.1816	-	$H0_{APP_{ARS}}$ not rejected
APP_{DOM}	0.2796	0.8569	0.4296	-	$H0_{APP_{DOM}}$ not rejected
APP_{LIK}	0.0024*	0.165	0.6368	0.6048 (large)	$H0_{APP_{LIK}}$ rejected in favor of YW
IMP_{PLS}	0.2008	0.6663	0.9793	-	$H0_{IMP_{PLS}}$ not rejected
IMP_{ARS}	0.6799	0.6881	0.5752	-	$H0_{IMP_{ARS}}$ not rejected
IMP_{DOM}	0.3449	0.5614	0.4672	-	$H0_{IMP_{DOM}}$ not rejected
IMP_{LIK}	0.0396*	0.1862	0.2703	0.4286 (medium)	$H0_{IMP_{LIK}}$ rejected in favor of YW
TES_{PLS}	0.0178*	0.65	0.7652	0.5 (large)	$H0_{IMP_{PLS}}$ rejected in favor of YW
TES_{ARS}	0.4147	0.4765	0.3406	-	$H0_{TES_{ARS}}$ not rejected
TES_{DOM}	0.6341	0.2564	0.4738	-	$H0_{TES_{DOM}}$ not rejected
TES_{LIK}	0.0504	0.1194	0.0547	-	$H0_{TES_{LIK}}$ not rejected

* P-value indicating a significant effect.

Further Analysis and Results. To better contextualize the baseline experiment, Romano *et al.* also assessed participants' development performance. To this end, they used a *time-fixed* strategy [2]. In particular, they defined an additional DV, named **STR**, which was computed as follows: (i) count the number of user stories each participant implemented within the fixed time frame (*i.e.*, three hours); then (ii) normalize the number of implemented user stories in [0, 100]—this is because the total number of user stories of MRA was different

¹ The descriptive statistics were used to determine if the difference was in favor of TDD or YW.

to that of BSK (*i.e.*, 11 vs. 13). It is ease to grasp that the higher the STR value is, the better the development performance of a given participant is. Romano *et al.* analyzed the effects of Approach, Object, and Approach:Object on STR by using ATS because the normality assumption to apply ANOVA [27] was not met. The results of ATS did not indicate a significant effect of Approach (p-value = 0.4765) on STR, namely the development approach seems not to influence the participants' development performance. The effects of Object (p-value = 0.2596), and Approach:Object (p-value = 0.0604) on STR were not significant.

Table 2. Summary of baseline and replicated experiments.

Characteristic	Baseline experiment	Replication
Participant type	III-year undergraduates in CS taking the SE course at the University of Basilicata	II-year undergraduates in CS taking the SE course at the University of Bari
Participant number	29	59
Site	University of Basilicata	University of Bari
RQs	RQ1, RQ2, RQ3	RQ1, RQ2, RQ3
Experimental objects	BSK, MRA	BSK, MRA
Experimental groups	TDD, YW	TDD, YW
Environment	Java, Eclipse, JUnit	Java, Eclipse, JUnit
Design	2 * 2 factorial	2 * 2 factorial
Assignment to groups and objects	15 participants assigned to TDD (7 BSK, 8 MRA), 14 participants assigned to YW (7 BSK, 7 MRA)	28 participants assigned to TDD (14 BSK, 14 MRA), 31 participants assigned to YW (16 BSK, 15 MRA)
IV	Approach, Object	Approach, Object
DV	APPPLS, APPARS, APPDOM, APPLIK, IMPPLS, IMPARS, IMPDOM, IMPLIK, TESPLS, TESARS, TESDOM, TESLIK	APPPLS, APPARS, APPDOM, APPLIK, IMPPLS, IMPARS, IMPDOM, IMPLIK, TESPLS, TESARS, TESDOM, TESLIK
Null hypotheses	H_{0DV}	H_{0DV}
Statistical inference method	ATS to analyze the effects of Approach, Object, and Approach:Object	ATS to analyze the effects of Approach, Object, and Approach:Object

4 Replicated Experiment

We conducted a replicated experiment to determine whether the results from the baseline experiment are still valid in a different site and with a larger number of participants sampled from a different population. Despite these differences, we designed and executed the replicated experiment as similarly as possible to the baseline experiment to determine, in case of inconsistent results with the baseline experiment, which factors could have caused those results. To this end, we used the replication package of the baseline experiment, which is available on the web² and includes experimental objects, analysis scripts, and raw data.

² <https://doi.org/10.6084/m9.figshare.9778019.v1>.

As shown in Table 2, the replicated experiment shares most of the characteristics of the baseline one. Therefore, in the following of this section, we limit ourselves to describe the replicated experiment in terms of participants, and design and execution. This is to say that RQs, artifacts, variables, hypotheses, and data analysis of the replication are the same as the baseline experiment; therefore, such information can be found in Sect. 3.

4.1 Participants

The participants in the replication were 59 second-year undergraduates in CS at the University of Bari who were taking the SE course. Participation was on a voluntary basis (*i.e.*, we did not pay the students for their participation). To encourage students to participate in the replication, we rewarded the participants with two bonus points in the final mark of the SE course (as had been done in the baseline experiment). The two bonus points were given regardless of the performance of the participants in the replication. Similarly to the baseline experiment, the participants were asked to fill in a pre-questionnaire. Based on the participants' answers, they had passed the exams of the Basic and Advanced Programming courses and had experience with C and Java programming. The participants were not knowledgeable in TDD.

4.2 Design and Execution

Based on the 2×2 factorial design used in the baseline experiment, the participants in the replication were randomly assigned to the experimental groups and objects: 28 participants were assigned to TDD—14 with BSK and 14 with MRA—; while 31 participants were assigned to YW—16 with BSK and 15 MRA.

All the participants in the replication attended face-to-face lessons on unit testing, JUnit, TL, and ITL. They also practiced unit testing with JUnit in a laboratory session. Later, the participants in the treatment group learned TDD and practiced it through two laboratory sessions and two homework assignments. The participants in the control group, who did not learn TDD, practiced TL and ITL through two laboratory sessions and two homework assignments. The material (*e.g.*, homework assignments) used to train the participants was the same as the baseline experiment, although the number of the homework assignments was different between the baseline and replicated experiments—*i.e.*, three vs. two. We were forced to give two homework assignments, rather than three, because the students could not carry out a third homework assignment during the training period due to deadlines that other courses requested in the same period. As so, we preferred not overloading students to avoid threat of dropouts from the experiment. We conducted the experiment in a single three-hour laboratory session in which the participants carried out the development task—*i.e.*, they tackled MRA or BSK—by using TDD or YW based on their experimental group. At the end of the development task, the participants self-assessed their affective reactions to the used development approach, as well as to implementing and testing code, by using SAM [3] complemented with the liking dimension [17].

5 Results

In Fig. 2, we summarize the values of the DVs (of the replicated experiment) by using diverging stacked bar plots. These plots show the frequencies of the DV values grouped by Approach. For each DV, the neutral judgment (*i.e.*, five) is displayed in grey; while negative judgments (*i.e.*, from one to four) and those positive (*i.e.*, from six to nine) are shown in shades of red and blue, respectively. The width of a colored bar (*e.g.*, the grey one) is proportional to the frequencies of the corresponding DV value (*e.g.*, five in the corresponding DV value for the grey bar). The interested reader can find the raw data on the web.³ The p-values ATS returned for each DV are reported in Table 3.

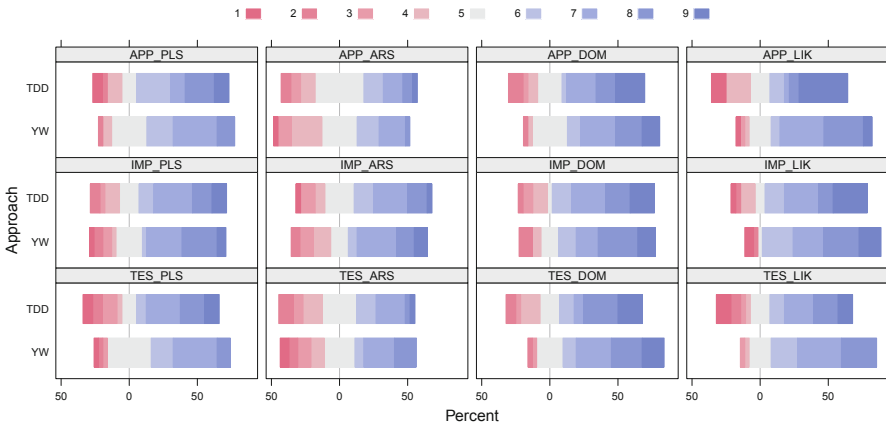


Fig. 2. Diverging stacked bar plots summarizing the DV values of the replication. (Colour figure online)

RQ1—Affective Reactions to the Development Approach. The plots in Fig. 2 (see the first row) do not show huge differences in the affective reactions to the used development approach, namely TDD or YW, in terms of pleasure (APP_{PLS}), arousal (APP_{ARS}), dominance (APP_{DOM}), and liking (APP_{LIK}). However, it seems that TDD has some negative frequencies more than YW as far as the dominance and liking dimensions are concerned. The results of ATS (see Table 3) indicate that there is no significant effect of Approach on the pleasure, arousal, dominance, and liking dimensions of the participants' affective reactions to the development approach. Accordingly, we cannot reject the corresponding null hypotheses. Finally, we did not find any significant effect of the interaction between Approach and Object, while the effect of Object is significant on the liking dimension (p-value = 0.0324). That is, the used experimental object significantly influenced the affective reactions of the participants to the development approach in terms of liking. However, the effect of the experimental object is consistent within both experimental groups as there is no significant interaction.

³ <https://doi.org/10.6084/m9.figshare.12085821.v1>.

Table 3. Results, from statistical inference, of the replication.

DV	IV			Outcome for H_{0DV}
	Approach	Object	Approach:Object	
APP _{PLS}	0.6937	0.0805	0.7001	$H_{0APP_{PLS}}$ not rejected
APP _{ARS}	0.6421	0.9018	0.2817	$H_{0APP_{ARS}}$ not rejected
APP _{DOM}	0.8295	0.1376	0.5235	$H_{0APP_{DOM}}$ not rejected
APP _{LIK}	0.9211	0.0324*	0.2571	$H_{0APP_{LIK}}$ not rejected
IMP _{PLS}	0.904	0.2849	0.4421	$H_{0IMP_{PLS}}$ not rejected
IMP _{ARS}	0.7781	0.9646	0.3198	$H_{0IMP_{ARS}}$ not rejected
IMP _{DOM}	0.9529	0.2389	0.9411	$H_{0IMP_{DOM}}$ not rejected
IMP _{LIK}	0.8048	0.1314	0.6618	$H_{0IMP_{LIK}}$ not rejected
TES _{PLS}	0.5722	0.3083	0.7749	$H_{0TES_{PLS}}$ not rejected
TES _{ARS}	0.7446	0.2281	0.4129	$H_{0TES_{ARS}}$ not rejected
TES _{DOM}	0.509	0.1079	0.9945	$H_{0TES_{DOM}}$ not rejected
TES _{LIK}	0.4588	0.3457	0.1566	$H_{0TES_{LIK}}$ not rejected

* P-value indicating a significant effect.

Answer to RQ1. We observed no significant difference in the affective reactions of novice developers to the used development approach, *i.e.*, TDD or YW.

RQ2—Affective Reactions to the Implementation Phase. As shown in Fig. 2, there is no huge difference between TDD and YW regarding pleasure (IMP_{PLS}), arousal (IMP_{ARS}), dominance (IMP_{DOM}), and liking (IMP_{LIK}) of the affective reactions to the implementation phase. We can also notice that, as for the liking dimension, TDD seems to have some negative frequencies more than YW. The results of ATS (see Table 3) do not show any significant effect of Approach on the four dimensions. Therefore, the corresponding null hypotheses cannot be rejected. The effects of Object and its interaction with Approach are not significant.

Answer to RQ2. With respect to the implementation phase, the results do not show a significant difference in the affective reactions of novice developers when they use TDD or YW.

RQ3—Affective Reactions to the Testing Phase. The plots in Fig. 2 show that the affective reactions of the control group to the testing phase in terms pleasure (TES_{PLS}), arousal (TES_{ARS}), dominance (TES_{DOM}), and liking (TES_{LIK}) are similar to the those of the treatment group. However, except for the arousal dimension, a slight trend in favor of YW can be observed since there are more negative frequencies for TDD as compared to YW. The results in Table 3 do not allow rejecting the null hypotheses. Finally, neither the effect of Object nor its interaction with Approach is significant.

Answer to RQ3. We did not observe a significant difference in the affective reactions of novice developers to the testing phase when they use TDD or YW.

Further Analysis Results. We used ATS to analyze STR because the normality assumption of ANOVA was not met (Shapiro-Wilk normality test p-value = 0.001). The results of ATS do not indicate a significant effect of Approach (p-value = 0.448) on STR, while the effect of Object (p-value < 0.001) was significant so suggesting that there was a difference in the development performance of the participants when dealing with BSK or MRA. However, the effect of the experimental object is consistent within both experimental groups since the interaction Approach:Object (p-value = 0.566) is not significant.

6 Discussion

Replications that do not draw the same conclusions as the baseline experiment can be viewed as successful, on a par with replications that come to the same conclusions as the baseline experiment [24]. Our replication falls into the former case since the outcomes of the replicated experiment do not fully confirm the outcomes of the baseline one. In particular, the baseline experiment found that participants seem to: (i) like TDD less as compared to YW; (ii) like less implementing code with TDD; and (iii) be less happy when testing code using TDD. The replication cannot support these findings because we did not observe any significant difference between TDD and YW. As for the other investigated constructs (*e.g.*, arousal due to the used development approach), the outcomes of the baseline experiment are confirmed by those of the replicated experiment (*i.e.*, the statistical conclusions are the same).

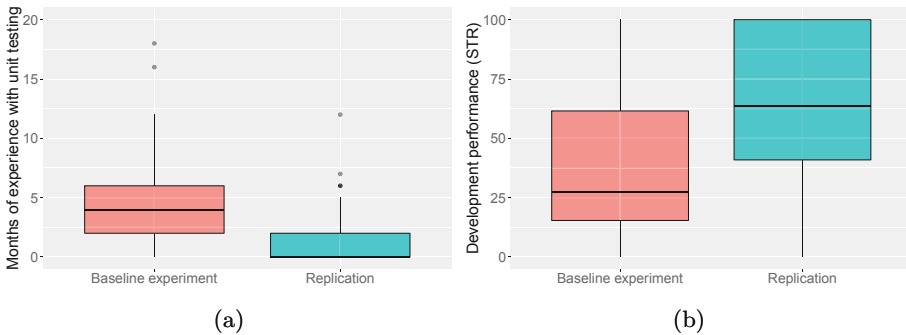


Fig. 3. Box-plots summarizing (a) months of experience with unit testing (at the beginning of the SE courses) of the participants and (b) development performance of the participants in the baseline and replicated experiments.

The question that now arises is why the replication fails to fully support the findings of the baseline one. We speculate that the inconsistent results between

the baseline and replicated experiments are due to the type of participants (third-year vs. second-year undergraduates in CS from two different universities), rather than their number (29 vs. 59). Although the number of participants in the baseline experiment was not so high and less than that of the participants in the replication, the magnitude (*i.e.*, Cliff's δ effect size) of the three significant effects [5], in the baseline experiment, was either medium or large. Such a magnitude makes us quite confident that the inconsistent results between the baseline and replicated experiments are not due to the number of participants. This is why we ascribe them to the type of participants. In particular, the participants in the baseline experiment were more experienced with unit testing than those in the replication, who mostly had no experience (see Fig. 3a). Since the participants in the baseline experiment did not know TDD (at the beginning of the SE course in which the experiment was run), they were therefore used to practice unit testing in a *test-last* manner. That is, they were used to write unit tests after they had written production code—in contrast to TDD, where unit tests are written before producing code. This is to say that the participants in the baseline experiment were probably more conservative and therefore less prone to change the order with which they usually wrote production and testing code. Accordingly, their affective reactions, due to TDD, were more negative. This postulation suggests two possible future research directions: (*i*) replicating the baseline experiment with more experienced developers to ascertain that the greater the experience with unit testing in a test-last manner, the more negative their affective reactions, due to TDD, are; and (*ii*) conducting an observational study with a cohort of developers to investigate if the affective reactions caused by TDD change over time. The above-mentioned postulation could be of interest to lecturers teaching unit testing. In particular, they could start teaching TDD as soon as possible to lessen/neutralize the negative affective reactions that TDD causes; after all, there is empirical evidence showing that, with time, TDD leads developers to write more unit tests [9].

Another characteristic of the participants that varies between the baseline and replicated experiments is the academic year of the CS program in which the participants were enrolled—*i.e.*, third year vs. second one. This implies that the participants in the baseline experiment have learned to code in Java a few months before than those in the replication. Nevertheless, the development performance was better in the replication than in the baseline experiment (see Fig. 3b). Therefore, we are quite confident that the academic year did not cause the inconsistent results between the baseline and replicated experiments. On the other hand, we cannot exclude that the worse development performance of the participants in the baseline experiment could have somehow amplified the differences in the affective reactions of the participants who practiced TDD or YW. After all, past work (*e.g.*, [11, 16]) has found that the affective states of developers are related to their performance in SE tasks, despite it is still unclear the role that TDD can play in such a relation. To better investigate this point, we suggest researchers to replicate the baseline experiment by introducing a change in the design, namely: allowing any participant to fulfil the development task (*i.e.*, no

fixed time), rather than giving any participant a fixed time frame to carry the development task. Such a design choice should allow isolating the effect that the development performance could have on the affective reactions of developers.

7 Threats to Validity

The replicated experiment inherits most of the threats to validity of the baseline one since, in the replicated experiment, we introduced few changes. We discuss the threats to validity according to the guidelines by Wohlin *et al.* [27].

Construct Validity. Threats concern the relation between theory and observation [27]. We measured each DV once by using a self-assessment instrument (*i.e.*, SAM). As so, in case of measurement bias, this might affect the obtained results (threat of *mono-method bias*). Although we did not disclose the research goals of our study to the participants, they might have guessed them and changed their behavior based on their guess (threat of *hypotheses guessing*). To mitigate a threat of *evaluation apprehension*, we informed the participants that they would get two bonus points on the final exam mark regardless their performance in the replication. There might be a threat of *restricted generalizability across constructs*. That is, TDD might have influenced some non-measured constructs.

Conclusion Validity. Threats concern issues that affect the ability to draw the correct conclusion [27]. We mitigated a threat of *random heterogeneity of participants* through two countermeasures: (*i*) we only involved students taking the SE course allowing us to have a sample of participants with similar background, skills, and experience; (*ii*) the participants underwent a training period to make them as more homogeneous as possible within the groups. A threat of *reliability of treatment implementation* might have occurred. For example, a few participants might have followed TDD more strictly than others, somehow influencing their affective reactions. To mitigate this threat, during the experiment, we reminded the participants to use the development approach we assigned them. Although SAM is one of the most reliable instruments for measuring affective reactions [19], there might be a threat of *reliability of measures* since the measures gathered by using SAM, as well as the liking scale, are subjective in nature.

Internal Validity. Threats are influences that can affect the IVs with respect to the causal relationship between treatment and outcome [27]. A *selection* threat might have affected our results since the participation in the study was on a voluntary basis and volunteers might be more motivated to carry out a development task than the whole population of developers. Another threat that might have affected our results is *resentful demoralization*, namely participants assigned to a less desirable treatment might not behave as they normally would. To mitigate a possible threat of *diffusion or treatments imitations*, we monitored the participants during the execution of the replication and alternated the participants dealing with BSK to those dealing with MRA.

External Validity. Threats to external validity concern the generalizability of results [27]. In the replication, we involved undergraduates in CS to reduce the

heterogeneity among the participants. This implies that generalizing the results to the population of professional developers might lead to a threat of *interaction of selection and treatment*. That is, while we mitigated a threat to conclusion validity like *random heterogeneity of participants*, we could not mitigate a threat to external validity. We prioritized a threat of *random heterogeneity of participants* to better determine, in case of different results between the baseline and replicated experiments, which factors might have caused such inconsistent results. However, it is worth mentioning that: (i) the use of students could be appropriate as suggested in the literature (e.g., [13,18,23]) and (ii) the development performance of the participants in the replication was better than that of the participants in the baseline experiment (see Fig. 3b). The use of BSK and MRA as experimental objects might represent a threat of *interaction of setting and treatment* despite they are commonly used as experimental objects in empirical studies on TDD (e.g., [9,10,23]). Moreover, both BSK and MRA can be fulfilled in a single three-hour laboratory session [9] so allowing better control over the participants.

8 Conclusion

We conducted a replicated experiment on the affective reactions of novice developers when applying TDD to implement software. With respect to the baseline experiment, we varied the experimental context and number of participants. The results from the replicated experiment do not fully confirm those of the baseline one. We speculate that the kind of developers can influence the affective reactions due to TDD. In particular, developers who have experience with unit testing in a test-last manner could have affective reactions, due to TDD, that are more negative than developers who have no/little experience with unit testing in a test-last manner. We also speculate that developers' performance in implementing software can influence the affective reactions of developers when applying TDD.

References

1. Beck, K.: Test-Driven Development: by Example. Addison-Wesley, Boston (2003)
2. Bergersen, G.R., Sjøberg, D.I.K., Dybå, T.: Construction and validation of an instrument for measuring programming skill. *IEEE Trans. Softw. Eng.* **40**(12), 1163–1184 (2014)
3. Bradley, M.M., Lang, P.J.: Measuring emotion: the self-assessment manikin and the semantic differential. *J. Behav. Ther. Exp. Psychiatry* **25**(1), 49–59 (1994)
4. Brunner, E., Dette, H., Munk, A.: Box-type approximations in nonparametric factorial designs. *J. Amer. Statist. Assoc.* **92**(440), 1494–1502 (1997)
5. Caivano, D.: Continuous software process improvement through statistical process control, pp. 288–293 (2005)
6. Carver, J.C., Juristo, N., Baldassarre, M.T., Vegas, S.: Replications of software engineering experiments. *Empir. Softw. Eng.* **19**(2), 267–276 (2014)

7. Cliff, N.: Ordinal Methods for Behavioral Data Analysis. Psychology Press, London (1996)
8. Erdogmus, H., Melnik, G., Jeffries, R.: Test-driven development. In: Encyclopedia of Software Engineering, pp. 1211–1229. Taylor & Francis (2010)
9. Fucci, D., et al.: A longitudinal cohort study on the retainment of test-driven development. In: Proceedings of International Symposium on Empirical Software Engineering and Measurement, pp. 18:1–18:10. ACM (2018)
10. Fucci, D., et al.: An external replication on the effects of test-driven development using a multi-site blind analysis approach. In: Proceedings of International Symposium on Empirical Software Engineering and Measurement, pp. 3:1–3:10. ACM (2016)
11. Graziotin, D., Wang, X., Abrahamsson, P.: Are happy developers more productive? In: Heidrich, J., Oivo, M., Jedlitschka, A., Baldassarre, M.T. (eds.) PROFES 2013. LNCS, vol. 7983, pp. 50–64. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39259-7_7
12. Herbon, A., Peter, C., Markert, L., Van Der Meer, E., Voskamp, J.: Emotion studies in HCI—a new approach. In: Proceedings of International Conference HCI (2005)
13. Höst, M., Regnell, B., Wohlin, C.: Using students as subjects—a comparative study of students and professionals in lead-time impact assessment. *Empir. Softw. Eng.* **5**(3), 201–214 (2000)
14. Kaptein, M.C., Nass, C., Markopoulos, P.: Powerful and consistent analysis of likert-type ratingscales. In: Proceedings of International Conference on Human Factors in Computing Systems, pp. 2391–2394. ACM (2010)
15. Karac, I., Turhan, B.: What do we (really) know about test-driven development? *IEEE Software* **35**(4), 81–85 (2018)
16. Khan, I.A., Brinkman, W.P., Hierons, R.M.: Do moods affect programmers' debug performance? *Cogn. Technol. Work* **13**(4), 245–258 (2011)
17. Koelstra, S., et al.: Deap: a database for emotion analysis using physiological signals. *IEEE Trans. Affect. Comput. Comput.* **3**(1), 18–31 (2012)
18. Lemos, O.A.L., Ferrari, F.C., Silveira, F.F., Garcia, A.: Development of auxiliary functions: should you be agile? an empirical assessment of pair programming and test-first programming. In: Proceedings of International Conference on Software Engineering, pp. 529–539. IEEE (2012)
19. Morris, J.D., Woo, C., Geason, J.A., Kim, J.: The power of affect: predicting intention. *J. Advert. Res.* **42**(3), 7–17 (2002)
20. Romano, J., Kromrey, J.D., Coraggio, J., Skowronek, J.: Appropriate statistics for ordinal level data: should we really be using t-test and Cohen'sd for evaluating group differences on the NSSE and other surveys? In: Annual Meeting of the Florida Association of Institutional Research, pp. 1–3 (2006)
21. Romano, S., Fucci, D., Baldassarre, M.T., Caivano, D., Scanniello, G.: An empirical assessment on affective reactions of novice developers when applying test-driven development. In: Franch, X., Männistö, T., Martínez-Fernández, S. (eds.) PROFES 2019. LNCS, vol. 11915, pp. 3–19. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-35333-9_1
22. Russell, J.A., Mehrabian, A.: Evidence for a three-factor theory of emotions. *J. Res. Personal.* **11**(3), 273–294 (1977)
23. Salman, I., Misirli, A.T., Juristo, N.: Are students representatives of professionals in software engineering experiments? In: Proceedings of International Conference on Software Engineering, vol. 1, pp. 666–676. IEEE (2015)
24. Shull, F.J., Carver, J.C., Vegas, S., Juristo, N.: The role of replications in empirical software engineering. *Empir. Softw. Eng.* **13**(2), 211–218 (2008)

25. da Silva, F., et al.: Replication of empirical studies in software engineering research: a systematic mapping study. *Empir. Softw. Eng.* **19**(3), 501–557 (2014)
26. Romano, S., Capece, N., Erra, U., Scanniello, G., Lanza, M.: The city metaphor in software visualization: feelings, emotions, and thinking. *Multimedia Tools Appl.* **78**(23), 33113–33149 (2019). <https://doi.org/10.1007/s11042-019-07748-1>
27. Wohlin, C., Runeson, P., Hst, M., Ohlsson, M.C., Regnell, B., Wessln, A.: *Experimentation in Software Engineering*. Springer, Heidelberg (2012)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

