Validation of Consumer and Research-Grade Activity Monitors Against
Polysomnography in Healthy Adults

Ria Jumabhoy, BA (Hons)¹, Svetlana Maskevich, BSc (Hons)¹, Pierre D.M Dao, BA
(Hons)/BSc,¹ Clare Anderson, Ph.D¹, Julie C. Stout, Ph.D,¹ & Sean P.A.
Drummond*, Ph.D¹

¹ Turner Institute for Brain and Mental Health, School of Psychological Sciences,
Monash University, Australia

*Corresponding Author:

Sean P.A. Drummond

Turner Institute for Brain and Mental Health

School of Psychological Sciences

Monash University

18 Innovation Walk

Clayton, VIC 3800, Australia

sean.drummond@monash.edu

ABSTRACT

Activity monitoring devices claiming to measure sleep have increased in popularity within research and consumer settings. However, validation studies on many of these devices are lacking. We examined twenty-nine healthy adults ($M_{\text{age}} = 32.4$, $SD_{\text{age}} = 11.6$) overnight using three activity monitors (Actiwatch Spectrum Pro standard and sensitive settings, Fitbit One, Jawbone UP2) and polysomnography. We evaluated agreement between each device and polysomnography on total sleep time (TST) and sleep efficiency (SE). Clinically meaningful limits of agreement were set at ± 30 minutes for TST and $\pm 5\%$ for SE. Each activity monitor overestimated TST and SE compared to polysomnography, and except for Actiwatch Spectrum Pro in sensitive setting, exceeded clinically meaningful agreement limits. Using Actiwatch Spectrum Pro sensitive as the gold standard, both consumer devices overestimated TST and SE, but their mean errors were generally not clinically meaningful (except Fitbit One's error for SE). Actiwatch Spectrum Pro in its sensitive setting may be a valid substitute for polysomnography when measuring TST and SE. Choice of consumer device as an actigraphy replacement will depend on context. Jawbone UP2 shows closer agreement with Actiwatch Spectrum Pro, though Fitbit One demonstrates less variation in TST and SE estimations.

Keywords:

Sleep
Sleep tracker
Actigraphy
Polysomnography
Validation

1. Introduction

Polysomnography (PSG) is the gold standard in sleep measurement. However, PSG is expensive to administer, time-consuming, impractical for long durations, and largely restricted to the laboratory.^{1,2} Activity monitors, or actigraphy devices, are typically worn on the wrist and estimate sleep and wake patterns through the use of accelerometers, which measure movement.^{1,3-5} Research-grade actigraphy devices are less expensive than PSG, but are still relatively costly and require training in administration and scoring.⁶

In healthy young adults and adolescents, actigraphs generally overestimate total sleep time (TST), wake after sleep onset⁷⁻¹⁰ (WASO), and sleep efficiency (SE: % of time in bed spent asleep), and underestimate sleep latency (SL).⁸⁻¹² Overestimations of TST typically range from 3.6 to 27.7 minutes, and overestimations of SE typically range from 0.7% to 5.0%. These errors are generally considered not to be clinically meaningful, with most authors placing clinically meaningful limits of agreement at ± 30 minutes for TST^{8,10,11,13,14} and $\pm 5\%$ for SE,^{8,10,13-15} though some studies show larger error in actigraphy.^{7,9,11,12,15} Importantly, each model of actigraph needs to be independently validated, as manufacturers often change software and/or hardware which can influence validity.

Recently, there has been a rise in popularity of consumer wearables tracking sleep-wake behaviour. Consumer sleep trackers also employ accelerometers, are significantly cheaper than research-grade actigraphy, and do not require training to administer or score. However, manufacturers do not reveal or allow users to alter algorithms used to measure sleep, and they do not easily provide data finer than all-

night measures. These represent inherent limitations to the current devices for research or clinical purposes. Nonetheless, the fact remains millions of individuals around the globe use these devices every day to assess their own sleep¹⁶. Thus, it is imperative the validity of these consumer devices be tested independently of the manufacturers.

Two of the more popular manufacturers of consumer sleep trackers are Fitbit and Jawbone. Validation studies on Fitbit, Fitbit Ultra, and Jawbone UP against PSG, have demonstrated these devices overestimate TST and SE, while underestimating WASO in healthy adults, adolescents, and youths.^{7,10,17,18} On TST, Fitbit overestimations have ranged from 41.0 to 67.1 minutes,^{15,19} while Jawbone overestimations ranged 9.1 to 26.6 minutes.^{10,17,18} On SE, Fitbit overestimations ranged 8% to 14.5%,^{15,19} whereas Jawbone overestimations ranged 1.8% to 1.9%.^{10,18} There have only been two studies reporting exceptions to these trends, with one showing Fitbit Flex overestimated TST and SE by <30 minutes and <5%, respectively,¹³ and another showing underestimations of both TST and SE for Fitbit Ultra in its sensitive setting.¹⁹

While PSG is the gold standard for validating sleep measurement devices and is the primary aim of this study, consumer devices should also be validated against research-grade actigraphy, since, in practice, they represent an alternative to actigraphy more so than to PSG. Three published studies^{15,19,20} have compared three different models of Fitbit (Fitbit, Fitbit Ultra, Fitbit Charge HR) to three different models of research-grade actigraphs (Actiwatch 64, Phillips Respironics Mini-Mitter Actiwatch-2, Actiwatch Spectrum Plus). Each study found Fitbit overestimated TST

(range 24–36 min) and SE (range 5.2–7%), relative to the actigraphs, although only two of the reports provided data sufficient to estimate the magnitude of error.^{15,19} To our knowledge, there have been no published studies validating any Jawbone model against research-grade actigraphy.

Overall, then, differences between research-grade actigraphy and PSG tend to remain within *a priori* defined clinically meaningful limits of agreement, although studies have reported both over and underestimations of actigraphy-based sleep parameters, relative to PSG. Findings on the Fitbit compared to PSG and actigraphy have been more consistent, with overestimations generally exceeding *a priori* limits for TST and SE. Research on Jawbone, on the other hand, remains less frequent and findings are less consistent.

To our knowledge, no studies have examined the validity of the Actiwatch Spectrum Pro (SP) model against PSG, compared the Jawbone UP2 to research-grade actigraphy, nor compared the accuracy of both the Fitbit and Jawbone (of any model) against both actigraphy and PSG within the same sample. The present study aims to fill these gaps in the literature with a sample of healthy adults. The current study focuses on TST and SE, as these variables incorporate SL and WASO into their calculations. The study had two primary aims: 1) examine differences in TST and SE between polysomnography and each of: the Actiwatch SP, Fitbit One, and Jawbone UP2; and 2) determine the sensitivity and specificity of Actiwatch SP, Fitbit One, and Jawbone UP2 compared to polysomnography. This study had two secondary aims: 1) examine differences in TST and SE between Actiwatch Spectrum Pro and each of: the

Fitbit One and Jawbone UP2; and 2) determine the sensitivity and specificity of Fitbit One and Jawbone UP2 compared to Actiwatch SP.

2 Materials and Methods

2.1 Participants

Participants were recruited online and from other projects running at the Monash University Sleep and Circadian Medicine Laboratory. All participants were studied on their adaptation nights prior to any experimental manipulation and to provide a stronger test of the devices, given the additional sleep disruption often seen on the first night in the laboratory. Ethics approval was gained from the Monash University Human Research Ethics Committee (CF16/422 – 2016000196) and the study was carried out in accordance with their guidelines. All participants read explanatory statements and gave informed consent.

Twenty-nine healthy participants (male=21; female=8) completed the overnight sleep study for this project ($M_{\text{age}} = 32.4$, $SD_{\text{age}} = 11.6$). One participant withdrew after providing consent for personal reasons.

Participants were excluded if they: 1) reported sleeping an average of <7 or >9 hours per night; 2) averaged >2 naps per day; 3) had a BMI >32; 4) could not speak English sufficiently to provide informed consent; 5) were currently involved in drug trials; 6) had current diagnoses of neurological, psychiatric or medical illnesses; 7) regularly used steroids, antipsychotics, or prescription or over-the-counter medication for sleep difficulties, asthma, epilepsy, or medication that affects the central nervous system; 8) had a history of loss of consciousness of >90 seconds or a blow to the head that resulted in migraines, seizures, cognitive difficulties, or the use of medication as a

result of the injury; 9) engaged in shift work within the last 3 months; 10) travelled to a different time zone within the last 3 months where 2 days had not passed for every time zone crossed; 11) drank an average of >3 standard alcoholic drinks in one sitting per week or an average of >2 standard drinks per day; 12) drank an average of >4 caffeinated drinks per day; 13) had used amphetamines, cocaine, methamphetamines, ecstasy or hallucinogens within the last month or marijuana within the last week; 14) were unwilling to abstain from cigarette smoking or nicotine replacement therapies during the overnight stay; or 15) obtained a global score of >5 on: the Pittsburgh Sleep Quality Index,²³ General Anxiety Disorder Scale-7,²⁴ or Patient Health Questionnaire-9,²⁵ and/or a score >10 on the Epworth Sleepiness Scale.²⁶ The extensive exclusion criteria were related to some of the studies that provided participants to this project. We adopted the most stringent of those criteria for all participants for consistency. The result is a very healthy set of normal controls.

2.2 Materials

2.2.1 Screening. Screening questionnaires included a demographics questionnaire, the Pittsburgh Sleep Quality Index, the General Anxiety Disorder-7 scale, the Epworth Sleepiness Scale, and the Patient Health Questionnaire-9 depression scale.

2.2.2 Polysomnography. PSG was applied according to the International 10/20 System of Electrode Placement. As recommended by the American Academy of Sleep Medicine (AASM),²⁷ measurements taken included frontal (F3, F4), central (C3, C4), and occipital (O1, O2) lobe electroencephalogram, electrocardiogram (ECG-L, ECG-R), chin electromyogram (M1, M2), leg electromyogram (EMG-L,

EMG-R), electrooculogram (EOG-L, EOG-R), respiration effort, oxygen saturation, and nasal and mouth airflow.

2.2.3 Activity Monitors. For research-grade actigraphy, we used the Actiwatch Spectrum Pro (SP) by Philips Respironics (2013 release). Epoch length was set at 1-minute intervals to match the default of the consumer devices. Consumer sleep trackers included the Fitbit One (2012 release) and the Jawbone UP2 (2015 release). Epoch length was pre-set at 1-minute intervals and this could not be altered.

2.3 Procedure

Potential participants contacted the researchers to gain information about the study, and to arrange a telephone screening and consent meeting. Upon passing the screening measures, they were invited for an overnight sleep study at the laboratory. Participants were required to remain in bed for a minimum of 6 hours for the devices to capture a realistic sleep period.²⁸ Watches were placed on participants' non-dominant wrist. The order of watches was counterbalanced across participants to account for the possibility that accuracy of devices may change according to the position on the arm. Devices began and finished recording simultaneously through manual synchronisation and the use of event markers to determine periods of analysis.

2.4 Data Extraction

Data from each device were analysed from time of lights off to time of lights on. PSGs were manually scored in one-minute epochs using Compumedics ProFusion PSG by viewing two 30-second epochs side-by-side, and treating them as one epoch. This is in line with previous studies using one-minute PSG epochs.^{10,17,19} Epochs

were scored in line with AASM rules as wake, N1, N2, or N3 (Version 2.1, 2014). For example, if two 30-second epochs together demonstrated 31 seconds of N1 and 29 seconds of N2, the one-minute epoch was scored as N1. Sleep stage data were then coded as either 'wake' or 'sleep', to be comparable to data extracted from watches. Minute-by-minute data from the Actiwatch SP were extracted using Philips Actiware 6.0 software. Actiwatch SP data were extracted using both standard settings (40 activity count wake threshold) and sensitive settings (20 activity count wake threshold) using the Cole-Kripke algorithm.³

Data from Fitbit One and Jawbone UP2 were not available as output files from consumer software. Thus, data for each epoch were manually extracted from summary graphs on their respective phone applications by scrolling across each graph and determining whether, in each minute, the watch recorded the participant as awake or asleep. Fitbit One indicates whether the participant is awake, restless, or asleep. As Fitbit One counts 'restless' as 'wake' in their calculations of sleep efficiency, we coded 'restless' as 'wake'.²⁹ Jawbone UP2 indicates whether, in each minute, the participant was awake, in light sleep, or sound sleep. Both light and sound sleep were treated together as 'sleep'. Fitbit One has a standard and a sensitive setting. However, unlike the Actiwatch SP, this refers to data collection mode, rather than data analysis mode (i.e., once data is collected in one mode, it cannot be retrospectively analysed in the other mode). Standard setting was used, because: a) it is the default; and b) Jawbone UP2 only has one setting (i.e., a standard setting).

2.5 Data Analysis

Alpha was set at 0.05, two-tailed, for all statistical analyses. Primary analyses included comparisons between gold standard PSG against Actiwatch SP, Fitbit One, and Jawbone UP2. Secondary analyses included comparisons between Actiwatch SP, and Fitbit One, and Jawbone UP2. No outliers were detected in any of the data. Jawbone UP2 failed to record for one participant, thus Jawbone UP2 comparisons are based on 28 participants. The reasons for recording failure are unclear and may be due to human or device error. Fitbit One data was not collected for two participants due to an inability to fit the watch on participants' arms alongside devices being used by other studies, thus Fitbit One comparisons are based on 27 participants.

2.5.1 Differences between devices. SPSS v22 was used to examine the statistical differences between devices. Assumption testing revealed violations of normality and normality of residuals. As such, the non-parametric alternative to a one-way ANOVA, Friedman's test, was employed to examine the differences in median TST and SE between devices.³⁰ Wilcoxon signed-rank tests were used to follow up significant findings.

2.5.4 Sensitivity and specificity. Using formulas outlined in Tilmanne et al.³⁶ sensitivity and specificity for each device was calculated. These calculations were conducted using true positives (TPs), true negatives (TNs), false positives (FPs), and false negatives (FNs). TPs occur when both the gold standard and the compared device has scored an epoch as sleep. TNs occur when both the gold standard and the compared device has scored an epoch as wake. FPs occur when the gold standard has scored an epoch as wake but the compared device has scored it as sleep. FNs occur when the gold standard has scored an epoch as sleep but the compared device has

scored it as wake. Sensitivity is calculated as $TP/(TP+FN)$ and refers to the percentage of epochs correctly identified by the compared device as sleep. Specificity is calculated as $TN/(TN+FP)$ and refers to the percentage of epochs correctly identified by the compared device as wake. Sensitivity and specificity were calculated individually for each participant, and Wilcoxon signed-rank tests were performed to assess significant differences in sensitivity and specificity values for each device. Effect sizes reported are r , the rank correlation between devices, and are calculated as the test statistic divided by the square root of the number of observations ($Z/\sqrt{N}$).³⁷

2.5.2 Interrater reliability. To assess agreement between devices beyond chance, Bennett's formula was used to calculate the prevalence and bias-adjusted kappa (PABAK).³¹ PABAK is suitable for assessing agreement in sleep research as it gives equal weight to sleep and wake epochs.¹¹ Using a confusion matrix, PABAK was calculated for the whole night across all participants, between each pair of devices. Landis and Koch's³² guidelines were used to interpret coefficient sizes.

2.5.3 Bland-Altman analyses. GraphPad Prism 6 software was used for Bland-Altman analyses. Bland-Altman plots assess the agreement between two devices by plotting the mean difference between devices against the average score, for each participant.^{33,34} While some authors³⁴ plot the difference score against the gold-standard score (PSG), rather than the average score, Bland and Altman³⁵ recommend against this method as it "will always appear to show a relationship between difference and magnitude when there is none".³⁵ Since proportional biases would be less meaningful with the use of this adjusted Bland-Altman plot, standard Bland-Altman plots were used for these analyses. Upper and lower limits of agreement are

defined as the mean difference between devices ± 1.96 standard deviations. Limits of agreement wider than *a priori* defined clinically meaningful limits demonstrate poor agreement between devices. Based on prior literature, we set *a priori* limits to ± 30 minutes for TST, and $\pm 5\%$ for SE.^{10,15}

3. Results

3.1 Primary Analyses Using PSG as the Gold Standard

3.1.1 Median Differences.

There were significant differences identified between devices for both TST, $\chi^2(4) = 57.96, p < .001$ and SE, $\chi^2(4) = 57.96, p < .001$. Compared to PSG, all devices overestimated TST and SE, with Actiwatch SP sensitive setting providing the closest estimates to PSG (Table 1). As the Actiwatch SP sensitive setting estimates were closer to PSG than the standard setting, all further actigraphy analyses were conducted solely on Actiwatch SP sensitive data. For parsimony, any future reference to Actiwatch SP will mean “Actiwatch SP in sensitive setting”.

3.1.2 Sensitivity and Specificity.

With PSG as the gold standard, Fitbit One and Jawbone UP2 were significantly more sensitive than Actiwatch SP ($p < .001, r = 0.83$ and $r = 0.72$ respectively), while Actiwatch SP was significantly more specific than Fitbit One ($p < .001, r = 0.84$) and Jawbone UP2 ($p = .001, r = 0.69$). See Table 2. No differences in sensitivity between Fitbit One and Jawbone UP2 were found ($p = .70$), but Jawbone UP2 was significantly more specific than Fitbit One ($p = .001, r = 0.61$).

3.1.3 Interrater Reliability.

On an epoch-by-epoch basis, agreement between PSG and Actiwatch SP was $\kappa = 0.76$ ($N = 13939$). Agreement between PSG and Fitbit One was $\kappa = 0.74$ ($N = 12977$), and agreement between PSG and Jawbone UP2 was $\kappa = 0.75$ ($N = 13458$).

3.1.4 Bland-Altman Method.

Compared to PSG, Fitbit One and Jawbone UP2's mean errors exceeded clinically meaningful limits of agreement for both TST and SE, while Actiwatch SP's mean errors remained within limits for both variables (Table 3, Figures 1-2). Upper and lower limits of agreement for all devices exceeded clinically meaningful limits. All devices demonstrated worse performance as sleep became worse (i.e. they showed proportional bias in their validity) for both TST and SE, except for Jawbone UP2 on SE (Table 3, Figure 2).

3.2 Secondary Analyses Using Actiwatch SP as the Gold Standard

3.2.1 Median Differences.

Compared to Actiwatch SP sensitive, both Fitbit One and Jawbone UP2 significantly overestimated both TST and SE ($p < .05$). See Table 1.

3.2.2 Sensitivity and Specificity.

With Actiwatch SP as the gold standard, no significant differences in sensitivity were found between Fitbit One and Jawbone UP2 ($p = .49$), but a significant difference in specificity was found between the two devices, with Jawbone UP2 showing greater specificity than Fitbit One ($p = .04$, $r = 0.39$). See Table 2.

3.2.3 Interrater Reliability.

On an epoch-by-epoch basis, agreement between Actiwatch SP and Fitbit One was $\kappa = 0.82$ ($N = 12957$) and agreement between Actiwatch SP and Jawbone UP2 was $\kappa = 0.79$ ($N = 13451$).

3.2.4 Bland-Altman Method.

When using Actiwatch SP in the sensitive setting as the gold standard, Fitbit One's mean error remained just under 30 minutes for TST and mean error for SE was just above the a priori cutoff of 5%. Limits of agreement generally were acceptable on the low end, but greatly exceed clinically acceptable levels on the high end (Table 4). Jawbone UP2's mean errors remained within clinically meaningful limits of agreement for both TST and SE, though both upper and lower limits of agreement exceeded those limits for both measures (Table 4). Fitbit One demonstrated a proportional bias in its validity against Actiwatch SP for both TST and SE, while Jawbone UP2 did not show proportional bias on either measure (Table 4, Figures 1-2).

3.3 Additional Analyses on Jawbone UP2

While Jawbone UP2's median values for TST and SE are closer to both PSG and Actiwatch SP than Fitbit One, the lower and upper limits of agreement for Jawbone UP2 are broader, suggesting more individual variability. This may be due to one instance in which Jawbone UP2 logged 203 minutes of consecutive wake during the night, when other devices recorded sleep. This led to underestimation of TST by 149 minutes and SE by 12.1%, compared to PSG, for this participant. All analyses were

rerun excluding this individual. The median for Jawbone UP2 remained closer to gold standards than Fitbit One, while maintaining broader lower and upper quartiles (Table S1 in supplemental material). Across the other analyses, the only qualitative difference was that Jawbone UP2 lost its specificity advantage over Fitbit One, when Actiwatch SP was used as the gold standard. (Tables S2-4 in supplemental material).

4. Discussion

This project aimed to validate the Actiwatch Spectrum Pro (SP), Fitbit One, and Jawbone UP2. To our knowledge, this is the first study to validate these specific models against PSG, and the first to compare both Fitbit One and Jawbone UP2 against research-grade actigraphy and PSG within the same sample. Based on median values, all three actigraphs statistically overestimated TST and SE compared to PSG, consistent with prior research (Table 1).^{9-12,14,15,17-19} Actiwatch SP in sensitive mode recorded median TST and SE closest to PSG, with mean errors on Bland-Altman plots for both TST and SE falling within *a priori* defined clinically meaningful limits of agreement (Table 3). On Bland-Altman plots, both Fitbit One and Jawbone UP2 made clinically meaningful errors in their estimations, when compared to PSG (Table 3). Compared to Actiwatch SP as the gold standard, both Fitbit One and Jawbone UP2 overestimated TST and SE based on median values, and Bland-Altman mean errors were generally not clinically meaningful (Table 4).

While mean error provides useful insight into the performance of sleep tracking devices, it does not provide information about individual variability around the mean, nor about whether sleep or wake is more likely to be misidentified. Such information is important for users to consider in the context of specific research and clinical

applications. Thus, we examined both limits of agreement (i.e., the mean difference between devices ± 1.96 standard deviations) and sensitivity/specificity. Although Actiwatch SP's mean error remained within clinically meaningful limits of agreement against PSG for TST and SE, and sensitivity was excellent, there was considerable individual variability in performance of the Actiwatch SP and specificity was <50%, failing to meet Sadeh's 60% benchmark²¹ (Tables 2 and 3). Poor specificity values suggest the device is missing periods of wake during the night. This could lead to inflated TST and SE values, which may lead researchers to misidentify individuals as normal sleepers, overestimate subjective-objective discrepancies in insomnia, or otherwise mischaracterise their samples. With all of that said, the values we report for Actiwatch SP are consistent with previous studies of other research actigraphy models.^{8,10,11,14,15} Thus, Actiwatch SP in sensitive setting, similar to other research-grade actigraphs, is a reasonable proxy for PSG, but tends to overscore sleep in the middle of the night, can show errors >60 minutes in TST and >15% in SE in specific individuals, and performance may become less accurate as sleep worsens.

Individual variability in performance of Fitbit One and Jawbone UP2, relative to PSG, was even greater than that seen in Actiwatch SP (Table 3). Specificity of both devices was also worse than Actiwatch SP and generally lower than previous research with consumer devices,^{10,15,17,19} potentially due to differences in models. The underestimation of wake during the night, and the resultant overestimation of both TST and SE, may be particularly problematic for consumers, as such errors can give a false impression regarding the individual's sleep health. In comparing the two consumer devices, relative to Actiwatch SP, Jawbone UP2 showed a wider range of individual variability than did Fitbit One, but better specificity (Tables 2 and 4).

While performance of Fitbit One declined with worsening sleep (i.e., shorter TST or lower SE), Jawbone UP2's performance did not vary with quantity or quality of sleep.

For research with healthy controls, the Actiwatch SP in its sensitive setting may be the most suitable alternative to PSG examined in this study. This is because both Fitbit One and Jawbone UP2 demonstrated poor specificity values, as well as limits of agreement and mean error values far exceeding clinically meaningful limits of agreement compared to PSG. The Actiwatch SP, on the other hand, maintained mean errors within clinically meaningful limits of agreement and demonstrated the highest level of specificity.

For researchers aiming to replace research-based actigraphs with a cost-efficient alternative, or consumers deciding between devices, it is not clear our results can recommend either device studied here. Jawbone UP2 did not show proportional bias against Actiwatch SP for TST and SE, and the percentage of cases exceeding *a priori* cut-offs is lower for Jawbone UP2 than for Fitbit One. On the other hand, at the individual level, Jawbone UP2 can show errors of >90 minutes in TST and >18% in SE, both more extreme than Fitbit One (Table 4). Thus, overall, Jawbone UP2 may show more accurate estimations of TST and SE than Fitbit One at the group level, while being more prone to large errors in individual cases. Regardless of relative strengths, though, both Fitbit One and Jawbone UP2 overestimated TST and SE by clinically meaningful amounts. Overestimations of TST have the potential to reach approximately 3 hours. This suggests, at an extreme, an individual may achieve 5 hours of sleep, but believe they achieved 7-8, potentially leading to unidentified chronic sleep restriction.²²

A few limitations of the current project are worth noting. Firstly, we had to score PSG in one-minute epochs, rather than 30-second epochs, to match output from both commercial devices. While not common outside the literature on device validation, using one-minute epochs is the most common strategy in validation studies using commercial activity monitors. Secondly, we did not assess whether differences between devices remained stable within participants over time. One key area for future research would be to replicate this study using a repeated measures design. If each device consistently overestimates or underestimates sleep variables to a similar degree within each participant, these devices can be used to monitor individual trends, and consumers can be more confident when they see changes in sleep as a consequence of changes in behaviour. Additionally, the present project used the standard setting for Fitbit One, as: 1) we could not alter sensitivity levels on Jawbone UP2; and 2) once Fitbit One data were collected in the standard setting, they could not be altered post-collection. Future research should examine the validity of Fitbit One in its sensitive setting. Importantly, we did not study a sleep-disordered population. Caution should be taken, especially with the proportional bias data, when extrapolating beyond the low ends of our TST and SE values. Finally, these models of consumer devices do not provide sleep stage data (other than “light” vs “sound” sleep in Jawbone UP2). Researchers and consumers alike will benefit from future validations of sleep stage estimates using newer models that allow for these analyses.

One final note: A limitation of all commercial sleep trackers is that none of the manufacturers make their scoring algorithms available to the public or researchers. This, in turns requires that each new device be independently validated before

consumers and researchers can truly have confidence in the data provided. Clearly, this is not ideal. Ultimately, it would be desirable for manufacturers to work with researchers by sharing relevant information on scoring algorithms, so that validation studies are no longer limited by 'black box' scoring procedures.

Acknowledgements

The authors thank Ms Charmaine Diep, Dr Suzanne Ftouni and Ms Jessica Manousakis for access to participants from their projects. We thank Ms Parisa Vidafar and Mr Christopher Andara for assistance with polysomnography scoring, Mr Brendan McLaren for assistance with the ethics application, and Mr Andrew Laing for his support during overnight shifts. We thank staff and students at the Monash University Sleep and Circadian Medicine Laboratory for their assistance throughout the project. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

1. Ancoli-Israel, S. *et al.* The role of actigraphy in the study of sleep and circadian rhythms. *Sleep* **26**, 342–392 (2003).
2. Sadeh, A. Iii. Sleep assessment methods. *Monogr. Soc. Res. Child Dev.* **80**, 33–48 (2015).
3. Cole, R. J., Kripke, D. F., Gruen, W., Mullaney, D. J. & Gillin, J. C. Automatic sleep/wake identification from wrist activity. *Sleep* **15**, 461–469 (1992).
4. Sadeh, A., Sharkey, K. M. & Carskadon, M. A. Activity-Based Sleep—Wake Identification: An Empirical Test of Methodological Issues. *Sleep* **17**, 201–207 (1994).
5. de Souza, L. *et al.* Further validation of actigraphy for sleep studies. *Sleep* **26**, 81–85 (2003).
6. Boyne, K., Sherry, D. D., Gallagher, P. R., Olsen, M. & Brooks, L. J. Accuracy of computer algorithms and the human eye in scoring actigraphy. *Sleep Breath.* **17**, 411–417 (2013).
7. Meltzer, L. J., Walsh, C. M., T aylor, J. & Westin, A. M. L. Direct comparison of two new actigraphs and polysomnography in children and adolescents. *Sleep* **35**, 159–166 (2012).
8. Paquet, J., Kawinska, A. & Carrier, J. Wake detection capacity of actigraphy during sleep. *Sleep* **30**, 1362–1369 (2007).
9. Tonetti, L., Pasquini, F., Fabbri, M., Belluzzi, M. & Natale, V. Comparison of two different actigraphs with polysomnography in healthy young subjects. *Chronobiol. Int.* **25**, 145–153 (2008).

10. Toon, E. *et al.* Comparison of Commercial Wrist-Based and Smartphone Accelerometers, Actigraphy, and PSG in a Clinical Cohort of Children and Adolescents. *J. Clin. Sleep Med.* **12**, 343–350 (2016).
11. Markwald, R. R., Bessman, S. C., Reini, S. A. & Drummond, S. P. A. Performance of a Portable Sleep Monitoring Device in Individuals with High Versus Low Sleep Efficiency. *J. Clin. Sleep Med.* **12**, 95–103 (2016).
12. Rupp, T. L. & Balkin, T. J. Comparison of Motionlogger Watch and Actiwatch actigraphs to polysomnography for sleep/wake estimation in healthy young adults. *Behav. Res. Methods* **43**, 1152–1160 (2011).
13. Mantua, J., Gravel, N. & Spencer, R. M. C. Reliability of Sleep Measures from Four Personal Health Monitoring Devices Compared to Research-Based Actigraphy and Polysomnography. *Sensors* **16**, (2016).
14. Weiss, A. R., Johnson, N. L., Berger, N. A. & Redline, S. Validity of activity-based devices to estimate sleep. *J. Clin. Sleep Med.* **6**, 336–342 (2010).
15. Montgomery-Downs, H. E., Insana, S. P. & Bond, J. A. Movement toward a novel activity monitoring device. *Sleep Breath.* **16**, 913–917 (2012).
16. Topic: Fitbit. *www.statista.com* Available at:
<https://www.statista.com/topics/2595/fitbit/>. (Accessed: 8th May 2018)
17. de Zambotti, M., Claudatos, S., Inkelis, S., Colrain, I. M. & Baker, F. C. Evaluation of a consumer fitness-tracking device to assess sleep in adults. *Chronobiol. Int.* **32**, 1024–1028 (2015).
18. de Zambotti, M., Baker, F. C. & Colrain, I. M. Validation of Sleep-Tracking Technology Compared with Polysomnography in Adolescents. *Sleep* **38**, 1461–1468 (2015).

19. Meltzer, L. J., Hiruma, L. S., Avis, K., Montgomery-Downs, H. & Valentin, J. Comparison of a Commercial Accelerometer with Polysomnography and Actigraphy in Children and Adolescents. *Sleep* **38**, 1323–1330 (2015).
20. Dickinson, D. L., Cazier, J. & Cech, T. A practical validation study of a commercial accelerometer using good and poor sleepers. *Health Psychology Open* **3**, 2055102916679012 (2016).
21. Sadeh, A. The role and validity of actigraphy in sleep medicine: an update. *Sleep Med. Rev.* **15**, 259–267 (2011).
22. Van Dongen, H. P. A., Maislin, G., Mullington, J. M. & Dinges, D. F. The cumulative cost of additional wakefulness: dose-response effects on neurobehavioral functions and sleep physiology from chronic sleep restriction and total sleep deprivation. *Sleep* **26**, 117–126 (2003).
23. Buysse, D. J., Reynolds, C. F., 3rd, Monk, T. H., Berman, S. R. & Kupfer, D. J. The Pittsburgh Sleep Quality Index: a new instrument for psychiatric practice and research. *Psychiatry Res.* **28**, 193–213 (1989).
24. Spitzer, R. L., Kroenke, K., Williams, J. B. W. & Löwe, B. A brief measure for assessing generalized anxiety disorder: the GAD-7. *Arch. Intern. Med.* **166**, 1092–1097 (2006).
25. Kroenke, K., Spitzer, R. L. & Williams, J. B. The PHQ-9: validity of a brief depression severity measure. *J. Gen. Intern. Med.* **16**, 606–613 (2001).
26. Johns, M. W. Reliability and factor analysis of the Epworth Sleepiness Scale. *Sleep* **15**, 376–381 (1992).
27. Kushida, C. A. *et al.* Practice parameters for the indications for polysomnography and related procedures: an update for 2005. *Sleep* **28**, 499–521 (2005).

28. Manousakis, J. The Sleep Habits of an Australian Adult Population: A Report on the 2015 Online Sleep Survey from the Sleep Health Foundation. (2015).
29. Fitbit. How do I track my sleep? (2016). Available at:
https://help.fitbit.com/articles/en_US/Help_article/1314.
30. Field, A. *Discovering Statistics Using IBM SPSS Statistics*. (SAGE, 2013).
31. Byrt, T., Bishop, J. & Carlin, J. B. Bias, prevalence and kappa. *J. Clin. Epidemiol.* **46**, 423–429 (1993).
32. Landis, J. R. & Koch, G. G. The measurement of observer agreement for categorical data. *Biometrics* **33**, 159–174 (1977).
33. Bland, J. M. & Altman, D. G. Statistical methods for assessing agreement between two methods of clinical measurement. *Lancet* **1**, 307–310 (1986).
34. Giavarina, D. Understanding Bland Altman analysis. *Biochem. Med.* **25**, 141–151 (2015).
35. Bland, J. M. & Altman, D. G. Comparing methods of measurement: why plotting difference against standard method is misleading. *Lancet* **346**, 1085–1087 (1995).
36. Tilmanne, J., Urbain, J., Kothare, M. V., Wouwer, A. V. & Kothare, S. V. Algorithms for sleep–wake identification using actigraphy: a comparative study and new results. *J. Sleep Res.* **18**, 85–98 (2009).
37. Pallant, J. F. *SPSS survival manual: A step-by-step guide to data analysis with SPSS*. (2007).

Author Contributions

RJ, SM, PDMD, JCS, and SPAD designed the study. RJ, SM, PDMD, CA, and SPAD contributed to recruitment of participants. RJ, SM, PDMD processed the data. RJ and SPAD conducted analyses. All authors contributed to data interpretation and preparation of the manuscript.

Data Availability

The datasets generated during and/or analysed during the current study are available from the corresponding author on reasonable request.

Declaration of Interest

All authors declare no relevant conflicts of interest.

Table 1.

Descriptive Statistics for TST and SE for PSG, Actiwatch standard, Actiwatch sensitive, Fitbit, and Jawbone

Variable	Device	Median	Interquartile range (lower-upper 25%)	Effect size relative to PSG (<i>r</i>)	Effect size relative to Actigraphy sensitive (<i>r</i>)
TST (minutes)	Polysomnography	438.0	386.0-452.0	-	-
	Actiwatch standard	458.0 ^a	443.5-466.5	0.77***	-
	Actiwatch sensitive	442.0 ^{ab}	417.5-453.5	0.43*	-
	Fitbit	469.0 ^{ac}	461.0-476.0	0.87***	0.86***
	Jawbone	464.5 ^{ac}	432.7-473.3	0.66**	0.69***
SE (%)	Polysomnography	88.6	81.2-93.9	-	-
	Actiwatch standard	95.8 ^a	92.9-97.0	0.77***	-
	Actiwatch sensitive	92.7 ^{ab}	88.2-94.2	0.42*	-
	Fitbit	97.5 ^{ac}	95.6-98.4	0.87***	0.86***
	Jawbone	95.8 ^{ac}	91.3-98.4	0.66***	0.69***

Note. TST = Total Sleep Time, SE = Sleep Efficiency. ^a = significantly different to PSG, ^b = significantly different to Actiwatch standard, ^c = significantly different to Actiwatch sensitive. * $p < .05$, ** $p < .01$, *** $p < .001$

Table 2.

Average Sensitivity and Specificity Values for Actiwatch, Fitbit, And Jawbone

Gold standard	Comparison Device	Sensitivity (%)	Specificity (%)
Polysomnography	Actiwatch	95.8	47.3
	sensitive		
	Fitbit	99.1 ^a	18.6 ^a
	Jawbone	97.2 ^a	32.2 ^{ab}
Actiwatch sensitive	Fitbit	91.6	20.2
	Jawbone	95.6	31.0 ^b

Note. ^a = significantly different from Actiwatch sensitive, ^b = significantly different from Fitbit. The specific *p*-values for each comparison are included in the text.

Table 3.

Bland-Altman Limits of Agreement for Activity Monitors and Polysomnography

Variable	Device	Mean error	Limits of agreement (lower– upper)	R ²	Cases exceeding <i>a priori</i> limits (%)
TST (minutes)	Actiwatch	23.8	-76.4, +124.0	0.40*	44.8
	sensitive				
	Fitbit	54.4	-65.1, +173.9	0.63*	55.2
	Jawbone	36.3	-90.5, +163.1	0.14*	58.6
SE (%)	Actiwatch	4.9	-15.6, +25.4	0.47*	51.7
	sensitive				
	Fitbit	11.2	-13.0, +35.4	0.90*	61.1
	Jawbone	7.6	-18.1, +33.3	0.14	61.1

Note. Significant R² demonstrates proportional bias. TST = Total sleep time, SE = Sleep efficiency. *A priori* limits were ± 30 minutes for TST and $\pm 5\%$ for SE. * $p < .05$, two-tailed

Table 4.

Bland-Altman Limits of Agreement for Consumer Devices and Actiwatch sensitive

Variable	Device	Mean error	Limits of agreement (lower– upper)	R ²	Cases exceeding <i>a priori</i> limits (%)
TST	Fitbit	29.2	-18.3, +76.8	0.28*	31.0
(minutes)	Jawbone	12.4	-72.0, +96.9	0.07	20.7
SE (%)	Fitbit	6.0	-3.5, +15.6	0.79*	44.8
	Jawbone	2.7	-14.4, +19.7	0.12	37.9

Note. Significant R² demonstrates proportional bias. TST = Total sleep time, SE = Sleep efficiency. *A priori* limits were ± 30 minutes for TST and $\pm 5\%$ for SE. * $p < .05$, two-tailed

Figure 1. Bland-Altman plots demonstrating upper and lower levels of agreement between Actiwatch and Jawbone for total sleep time. Solid line represents the mean error; dotted lines represent the upper and lower limits of agreement; trend line demonstrates proportional bias. Values above and below 0 demonstrate overestimation and underestimation of the device against the gold standard respectively. TST = Total sleep time.

Figure 2. Bland-Altman plots demonstrating upper and lower levels of agreement between Actiwatch and Jawbone for sleep efficiency. Solid line represents the mean error; dotted lines represent the upper and lower limits of agreement; trend line demonstrates proportional bias. Values above and below 0 demonstrate overestimation and underestimation of the device against the gold standard respectively. SE = Sleep efficiency.

Table S1.

Recalculated Descriptive Statistics for TST and SE for PSG, Actiwatch standard, Actiwatch sensitive, Fitbit, and Jawbone

Variable	Device	Median	Interquartile range (lower- upper 25%)
TST (minutes)	Polysomnography	433.0	383.8-452.8
	Actiwatch standard	457.5 ^a	443.3-465.0
	Actiwatch sensitive	441.0 ^a	416.8-451.8
	Fitbit	468.5 ^{abc}	459.3-475.3
	Jawbone	465.0 ^{ac}	435.0-474.0
SE (%)	Polysomnography	88.6	80.9-93.9
	Actiwatch standard	95.5 ^a	92.8-97.0
	Actiwatch sensitive	92.4 ^a	87.8-94.0
	Fitbit	97.5 ^{abc}	95.5-98.4
	Jawbone	96.2 ^{ac}	91.7-98.5

Note. TST = Total Sleep Time, SE = Sleep Efficiency. ^a = significantly different to PSG (all $p < .05$), ^b = significantly different to Actiwatch standard (all $p < .01$), ^c = significantly different to Actiwatch sensitive (all $p < .001$).

Table S2.

Recalculated Average Sensitivity and Specificity Values for Actiwatch, Fitbit, And Jawbone

Gold standard	Comparison Device	Sensitivity (%)	Specificity (%)
Polysomnography	Actiwatch sensitive	95.7	48.3
	Fitbit	99.0 ^a	18.6 ^a
	Jawbone	98.5 ^a	31.3 ^{ab}
Actiwatch sensitive	Fitbit	91.4	20.0
	Jawbone	97.0	29.4

Note. ^a = significantly different from Actiwatch sensitive (all $p < .05$), ^b = significantly different from Fitbit (all $p < .01$).

Table S3.

Recalculated Bland-Altman Limits of Agreement for Activity Monitors and Polysomnography

Variable	Device	Mean error	Limits of agreement (lower-upper)	R ²	Cases exceeding <i>a priori</i> limits (%)
TST (minutes)	Actiwatch sensitive	23.1	-76.9, +125.0	0.42*	42.9
	Fitbit	54.6	-67.2, +176.5	0.64*	53.6
	Jawbone	43.2	-63.8, +150.1	0.43*	57.1
SE (%)	Actiwatch sensitive	4.8	-16.1, +25.6	0.48*	50.0
	Fitbit	11.3	-13.4, +35.9	0.90*	60.7
	Jawbone	9.0	-12.8, +30.7	0.48*	60.7

Note. Significant R² demonstrates proportional bias. TST = Total sleep time, SE = Sleep efficiency. *A priori* limits were ± 30 minutes for TST and $\pm 5\%$ for SE. * $p < .05$, two-tailed

Table S4.

Recalculated Bland-Altman Limits of Agreement for Consumer Devices and Actiwatch sensitive

Variable	Device	Mean error	Limits of agreement (lower-upper)	R ²	Cases exceeding <i>a priori</i> limits (%)
TST (minutes)	Fitbit	30.1	-17.6, +77.7	0.26*	32.1
	Jawbone	20.0	-12.7, +52.6	0.001	17.9
SE (%)	Fitbit	6.2	-3.4, +15.8	0.79*	46.4
	Jawbone	4.2	-2.6, +10.9	0.0002	35.7

Note. Significant R² demonstrates proportional bias. TST = Total sleep time, SE = Sleep efficiency. *A priori* limits were ± 30 minutes for TST and $\pm 5\%$ for SE.