# The average shot length and the ecological fallacy in Film Studies

Nick Redfern
Independent scholar
Leeds, United Kingdom
E-mail: nickredfernres@outlook.com
ORCID: 0000-0002-7821-2404

**Abstract:**

The ecological fallacy is an error in the interpretation of statistical data when one ascribes to an individual the qualities of the group to which that individual belongs. In this article I show that using the average shot length (ASL) leads film researchers to commit the ecological fallacy due to the skewed nature of shot length distributions. The ecological fallacy can be avoided when analysing motion picture shot length data using methods that assess the stochastic equality of films' shot lengths.

**Keywords:** Computational film analysis, film editing, ecological fallacy, average shot length

For nearly 50 years (Salt 1974), film scholars have employed the average shot length (ASL) as a statistic of film style to describe and compare the editing of motion pictures. The ASL is the mean duration of the shots in a film (where a shot is a continuous sequence of frames) and is calculated as the sum of the durations of the shots in a film divided by the number of shots, though many scholars do not, in fact, measure the duration of individual shots in a film and simply divide the running time of a film by the number of shots to arrive at the same result. Conventionally, the ASL is interpreted as a measure of the cutting rate of a film, and when comparing the editing of two films, to answer the question 'do the shots in film *X* tend to have longer duration than the shots in film *Y*,' the difference between their ASLs is understood to be an estimate of the size of the difference in their editing style.

In this article I demonstrate that the use of the average shot length (ASL) as a statistic of film style leads researchers to commit an ecological fallacy by making inferences about a low-level variable (the duration of shots in motion pictures) based on a statistic (the ASL) that aggregates shot length data at a higher level (a film). I show that film scholars can avoid the ecological fallacy by assessing the stochastic equality of shot lengths when comparing editing data from motion pictures.

## The ecological fallacy

The ecological fallacy is an error in the interpretation of statistical data that arises from ascribing to an individual the characteristics of the group to which they belong (Brewer and Venaik 2014). Wu (2007, 123) describes ecological fallacies as 'a problem of disaggregation (or downscaling) in which inferences about a lower level are made from knowledge of an upper level.' Consequently, 'one commits an ecological fallacy if one assumes that relationships observed at an aggregated level imply that the same relationships exist at the individual level' (Jargowsky 2021, 48). There are different versions of the ecological fallacy relating to (1) confusion between ecological and individual correlations; (2) confusions between group averages and total averages; (3) Simpson's paradox, in which correlations are reversed under aggregation; and (4) the confusion between higher averages and higher likelihood.

Here we are concerned with the fourth type of ecological fallacy, which Brewer and Venaik, (2014, 1067) dub the 'averaging ecological fallacy,' with high-level aggregate data in the form of averages used to make statements about individuals at a lower level producing a conflict between the unit of observation (i.e., the object for which data is collected) and the unit of analysis (the object for which data is analysed and conclusions made) (Sedgwick 2015). For example, just because a student is a member of a class that has a high average score on a test does not mean we can assume that that student scored highly on the test. The student may be a low scoring member in an otherwise high scoring class; or she may be a mediocre performer in a class that includes a small number of outstanding students who scored exceptionally highly. In this example, the unit of observation at which measurements are made are the individual students but the unit of analysis at which inferences are drawn is the collective (i.e., the class). It is an ecological fallacy to assume that inferences made at the level of the class can be applied to the individual students because the average attribute of a class does not necessarily apply to all members of that class.

A conflict between the level of the unit of observation and the level of the unit of analysis is evident in the use of the ASL in Film Studies. For example, Adelheid Heftberger (2018, 148) states that

> a high ASL thus means that on average, the film contains longer shots with fewer cuts.

Similarly, Kaire Maimets-Volt (2013, 67n19) states that

> a long ASL means the film uses, on average, longer shots and fewer cuts.

It is clear from these interpretations of the average shot length that a high-level aggregate (the ASL) at the level of the unit of analysis (the film) is used to make inferences at the lower level of the unit of observation (the individual shots in a film), where a 'high' or 'longer' ASL means 'longer shots.' As we shall see, this will lead film scholars into error when they conclude that the shots in a film tend to have greater duration than shots in a second film because that film has a higher ASL.

***Une femme mariée*** **(1964) and *Week-end* (1967)**

To illustrate how the ASL leads us to commit an ecological fallacy I compare the shot lengths for two films directed by Jean-Luc Godard – *Une femme mariée* (1964) and *Week-end* (1967) – using data from the Cinemetrics database (Salt 2007a, 2007b).

Table 1 presents a statistical summary of the shot lengths in the two films, including the number of shots, the ASL, and the five-number summary comprising the minimum shot length, lower quartile, median shot length, upper quartile, and maximum shot length for each film. The lower and upper quartiles cut off the lower and upper 25% of a data set, while the median divides the range of the data into two equal parts.

Table 1. Summary of shot length data in two Jean-Luc Godard films.

| Statistics | *Une femme mariée* | *Week-end* |
|---|---|---|
| **Shots** | 211 | 231 |
| Running time (s) | 4750.3 | 6170.3 |
| Average shot length (s) | 22.5 | 26.7 |
| Minimum (s) | 0.8 | 0.2 |
| Lower quartile (s) | 6.0 | 2.0 |
| Median (s) | 13.0 | 6.7 |
| Upper quartile(s) | 33.7 | 26.0 |
| Maximum (s) | 119.5 | 474.0 |

From Table 1 we see that *Une femme mariée* has an ASL of 22.5 seconds and that *Week-end* has an ASL of 26.7 seconds. Going by the conventional use of the ASL in Film Studies, this would be interpreted as meaning that *Week-end* contains 'longer shots with fewer cuts' than *Une femme mariée*. However, this is clearly not true. Looking at the number of shots in each film in Table 1 shows that *Week-end* has 231 shots and that *Une femme mariée* has 211 shots. The conventional interpretation that a higher ASL means that a film will have fewer cuts is obviously false. In fact, it is not possible to arrive at any conclusion about which film will have fewer cuts based on the ASL.

From Table 1 we can also see that, although the ASL for *Week-end* is greater than that of *Une femme mariée*, the minimum, lower quartile, median, and upper quartile shot lengths of *Une femme mariée* are greater than those of *Week-end*, while the maximum shot length in *Week-end* is much greater than that of *Une femme mariée*. This is a clear sign that the ASL is leading us towards an erroneous conclusion because, as the five-number summary indicates, the majority of shots in *Une femme mariée* have longer duration than those of *Week-end*, while the latter film has some shots with very long duration.

It is easier to see the difference between the shot lengths of the two films if we visualise the data. Figure 1 plots the empirical cumulative distribution functions (ECDF) of the shot lengths in *Une femme mariée* and *Week-end*. The ECDF plots the proportion of shots in a film that have duration less than or equal to a specified value (X):

$$F(x) = \frac{\#(x \leq X)}{N},$$

where #(x ≤ X ) is the number of shots (x) with duration less than or equal to X and N is the total number of shots in a film. Visually comparing the shot lengths in two films is straightforward because the ECDF of a film that has a greater proportion of shots less than or equal to some value X will lie to the right of the other film.

From Figure 1 we see that the ECDF of *Une femme mariée* lies to the right of that of *Week-end* for shot lengths up to 43.8 seconds, indicating that it tends to have longer shots up to this value, but that the ECDF of *Week-end* lies to the right of that of *Une femme mariée* for shot lengths greater than 43.8 seconds, which accounts for 15% of the shots in these films. We can therefore conclude that shot lengths in *Une femme mariée* tend to be greater than those of *Week-end*, though the longest shots in the latter film have much greater duration. As we saw when looking at the five-number summary, this is clearly very different picture of the style of these two films provided by comparing their respective ASLs.

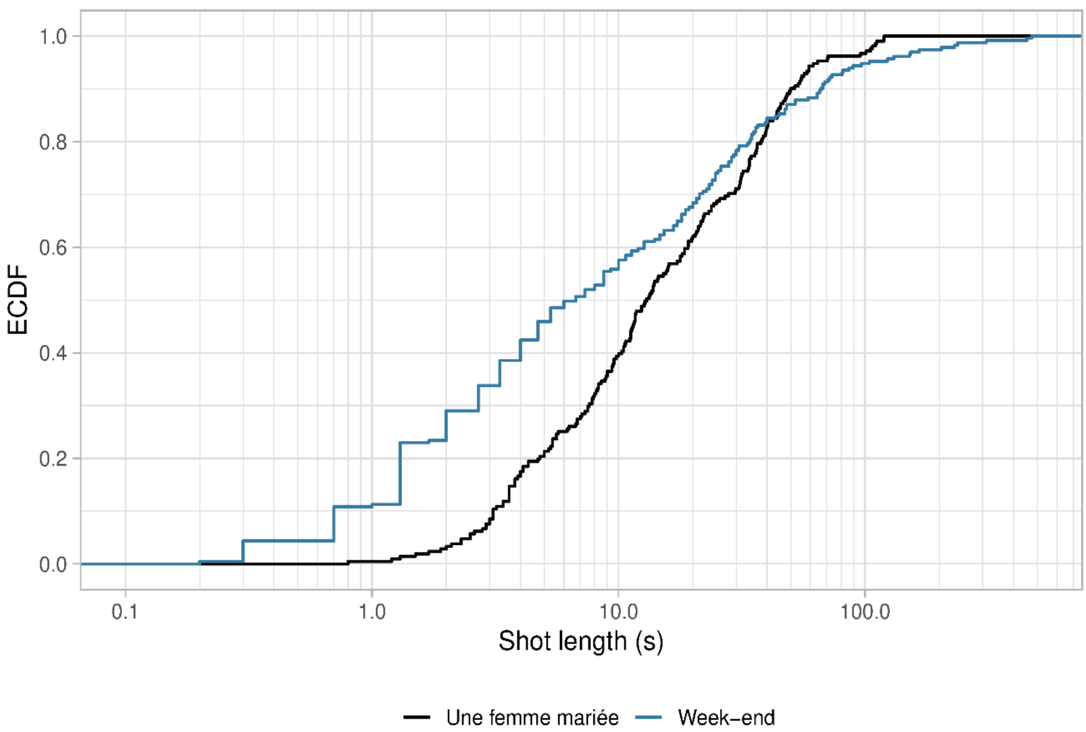


Figure 1. Empirical cumulative distribution functions of shot lengths in Jean-Luc Godard's *Une femme mariée* (1964) and *Week-end* (1967).

## Why does the ecological fallacy occur for shot length data?

As we have seen above, summary statistics can 'sometimes conceal or even misrepresent what may be the most informative aspects of the data' (Hartwig and Dearing 1979, 9), and film scholars place too much trust in the ASL as a summary statistic of film style. The ecological fallacy occurs when we uncritically apply a group-level aggregate (i.e., a film's ASL) to the all the individuals in that group (i.e., the shots in that film), ignoring the variation that exists within that group.

To understand why using the ASL leads us to commit an ecological fallacy, we therefore need to grasp the concept of a *shot length distribution*. A distribution is a representation of the variation in a data set and is the lens through which we investigate the patterns in the data (Wild 2006), attending to attributes including the centre (where is the mass of the data located?), modality (how many peaks are there?), spread (how much variability is there?), skewness (is the distribution symmetrical?), peakedness (are the peaks flat and broad or tall and pointed?), tailedness (how much data is located in the tails relative to the centre?), and deviations from the overall pattern of the data (e.g., the presence of outliers).

The data set of shot lengths collected from a film exhibits variation: *the shots in a film have different durations, they vary, and a shot length distribution is a representation of the variation of the shot lengths in a film*. It may seem unnecessary to make such a simple observation but few film scholars applying quantitative methods to editing in the cinema employ the concept of a shot length distribution. One reason for this is that many films scholars do not collect any shot length data and calculate the ASL by dividing the running time of the film by the number of shots. They cannot understand the variation in their shot length data because *they have no data*. But even when scholars do collect shot length data, they typically report only the ASL and fail to attend to the variation in the data. Consequently, film scholars discuss the shot lengths of films in terms of their average as if they exhibited no internal variation (Wenkert 1961) and do not address key features of the distribution of their data set. Without knowledge of the distribution of shot lengths in two films (i.e., information about the unit of observation) concluding that a film with a higher ASL tends to have shots of longer duration raises the dangers of committing an ecological fallacy because it is possible that the likelihood of randomly selecting a shot with greater duration in the film with the lower ASL is not reflected by the average shot length.

The averaging ecological fallacy occurs because the distributions of the shot lengths in motion pictures are, heavily and differentially, *positively skewed*. That is, they are not symmetrical with most of the data points concentrated in the left tail of the distribution with a long right tail. Figure 2 plots the distribution of shot lengths in *Une femme mariée* and *Week-end*. We can clearly see that the shot length distributions of both films are positively skewed. *Week-end* has a small number of shots that are over 150 seconds in duration that are much longer than the longest shots in *Une femme mariée*. Because the ASL is based on the arithmetic mean, it is highly sensitive to shots with very long duration which cause the value of the ASL to be pulled in the direction of the outlying data points. This is also evident in Table 1, which shows that the ASL in *Week-end* (26.7 seconds) is greater than even the upper quartile (26.0 seconds).

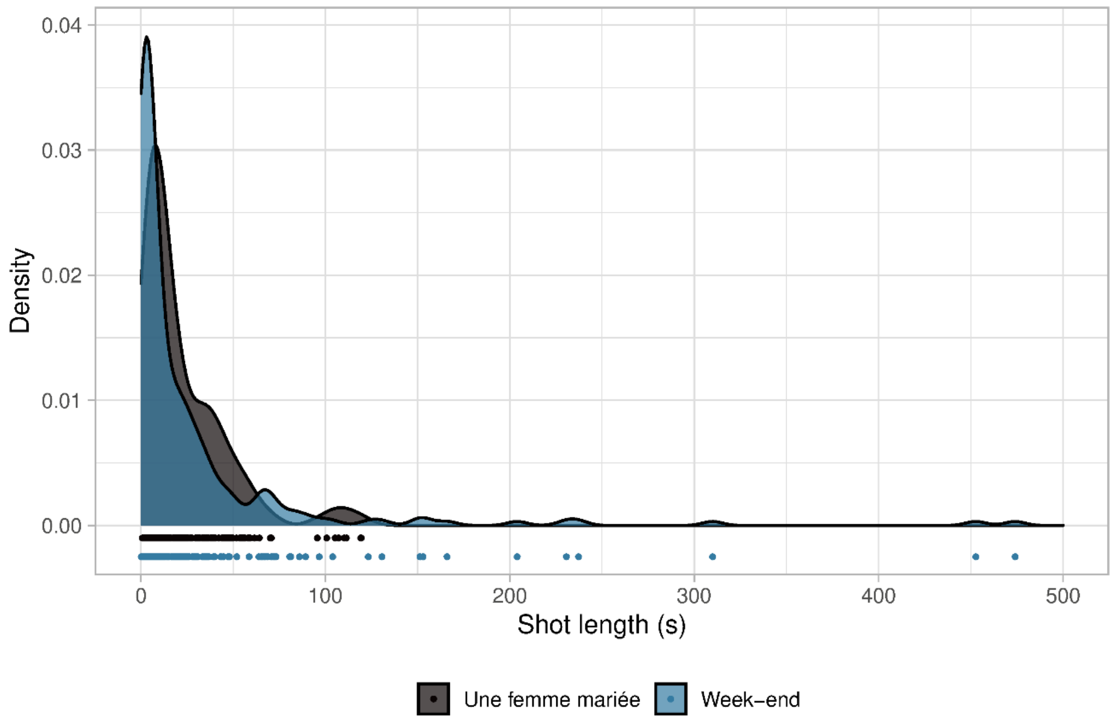


Figure 2. The distributions of shot lengths in Jean-Luc Godard's *Une femme mariée* (1964) and *Week-end* (1967) are positively skewed with the mass of the data concentrated to the left of the scale and long right tails.

The great variability in duration of shots within a film has important implications for the significance attached to the ASL. The ASLs of *Une femme mariée* and *Week-end* do not capture anything meaningful about the distribution of the shot lengths in these films. They do not locate the centre of the distributions. They tell us nothing about the spread of the data. They contain no information about the symmetry or tailedness of the distributions. The ASL is a vacuous statistic of film editing that forces us to make inferences about the individual shots in a film based on an aggregate at a higher level of data, creating a conflict between the unit of observation and the unit of analysis.

**The stochastic equality of shot lengths**

We can avoid committing the ecological fallacy when comparing the shot length data of two films by applying methods based on evaluating the stochastic equality of motion picture shot length data. This means we will analyse the shot lengths of our films directly and will not employ summary measures such as the ASL.

We can determine the likelihood that the length of a shot in *Week-end* is greater than or equal to the length of a shot in *Une femme mariée* by directly comparing the length of each shot in one film with length of each shot in the other to calculate Vargha and Delaney's *A* statistic (Vargha and Delaney 2000). The *A* statistic is equal to the number times the length of a shot in *Week-end* is greater than the length of a shot in *Une femme mariée* (#(WE > UFM)) divided by the product of the number of shots in *Une femme mariée* ($N_{UFM}$) and the number of shots in *Week-end* ($N_{WE}$) plus half the number of times a shot in *Une femme mariée* and a shot in *Week-end* have equal lengths (#(WE = UFM)) divided by the product of the number of shots in each film:

$$A = \frac{\#(WE > UFM)}{N_{WE} \times N_{UFM}} + 0.5 \frac{\#(WE = UFM)}{N_{WE} \times N_{UFM}}.$$

If $A = 0.5$ then there is no tendency for shots in one film to be longer: i.e., shots in *Week-end* as equally as likely to have longer duration than shots in *Une femme mariée* as they are to have shorter duration. A value of $A$ greater than 0.5 tells us that shot lengths in *Week-end* have a higher likelihood of having longer duration than shot lengths in *Une femme mariée*, while a value of $A$ less than 0.5 will indicate that shot lengths in *Week-end* have a lower likelihood of having longer shot lengths.

The $A$ statistic for *Une femme mariée* and *Week-end* is 0.38. This means that if we select at random a shot from *Une femme mariée* and we select at random a shot from *Week-end*, the probability that the latter will have longer duration that the former is 0.38. In other words, shots in *Week-end* have a lower likelihood of having longer duration than shots in *Une femme mariée* even though it has the higher ASL.

A limitation of the $A$ statistic is that it is not intuitive because stochastic equality occurs when $A = 0.5$. We can apply a simple linear transformation to obtain Cliff's $d$ statistic (Cliff 1993): $d = 2A - 1$. Cliff's $d$ statistic is equal to proportion of shots *Week-end* that have greater duration than shots in *Une femme mariée* minus the proportion of shots in *Week-end* with shorter duration than shots in *Une femme mariée*:

$$d = \frac{\#(WE > UFM)}{N_{WE} \times N_{UFM}} - \frac{\#(WE < UFM)}{N_{WE} \times N_{UFM}}.$$

$d$ is more intuitive to understand than A because stochastic equality occurs when $d = 0$. The sign of the d statistic will then tell us the direction of the difference. For *Une femme mariée* and *Week-end*, $d = -0.24$, indicating that shots in *Week-end* tend to have shorter duration than those in *Une femme mariée*.

The conventional interpretation of the ASL is that the difference between two films' ASLs is an estimate of the difference in their styles. Subtracting the ASL of *Une femme mariée* from the ASL of *Week-end* gives us a difference of 26.7 – 22.5 = 4.2 seconds, indicating that, on average, shots in *Week-end* are much longer in duration than those in *Une femme mariée*. However, we know that this cannot be correct because shots in *Week-end* tend to have shorter duration than those in *Une femme mariée*.

We can estimate the size of the difference between the shot lengths in these two films by calculating the Hodges-Lehmann median difference (Hodges and Lehmann, 1963), which is the median of all the pairwise differences produced by subtracting the length of every shot in *Une femme mariée* ($UFM_i$) from the length of every shot in *Week-end* ($WE_j$):

$$HL\Delta = median\left(WE_j - UFM_i\right).$$

For these two films, the Hodges-Lehmann median difference is −3.9 seconds. The sign of HLΔ is the opposite of difference between the ASLs and again tells us that shots in *Week-end* tend to be shorter than those in *Une femme mariée*, which aligns with the fact that shots in the former have a lower likelihood being of longer duration. Note that HLΔ and Cliff's *d* have the same sign.

A key advantage of using *A*, *d*, and HLΔ is that they allow us to make inferences about the duration of shots by analysing the duration of shots. Unlike the ASL, which forces us to make inferences about shot lengths at the aggregated level of the film, the units of observation and analysis are identical. Consequently, these methods allow us to answer the question 'do the shots in film *X* tend to be greater than the shots in film *Y*?' and to estimate the size of any difference.

**Conclusion**

The average shot length is considered to be a straightforward statistic of film style, a 'rather obvious concept' (Salt 1992, 146), deployed routinely by film scholars with a straightforward interpretation; but contrary to the definition of the average shot length used in Film Studies, a high ASL does *not* mean that, on average, a film contains longer shots with fewer cuts.

As the example of *Une femme mariée* and *Week-end* demonstrates, had we used the ASL as a statistic of film style we would have committed an ecological fallacy when interpreting the shot length data of these films, incorrectly concluding that shots in *Week-end* have longer duration than those in *Une femme mariée* when in fact the opposite is the case: *Week-end* has the higher ASL but the likelihood that a shot will have greater duration is higher for *Une femme mariée*. If our goal is to answer the question, 'do the shots in film *X* tend to be greater than the shots in film *Y*?,' then we must recognise that the ASL cannot answer that question.

Using statistics such as the five-number summary that describe shot length data, visualising the distribution of shot lengths in films, and using statistics such as *A*, *d*, and HLΔ to compare the style of films we are able correctly identify differences in style between films because when we make statements about shot lengths they actually will be statements about shot lengths. We will avoid committing the ecological fallacy because we will not be forced to make inferences about a lower-level variable based on aggregates at an upper level and we will not ascribe to individual shots the characteristics of the film to which they belong. We will be able to answer the questions about film style we wish to ask.

## References

Brewer, Paul, and Sunil Venaik. “The ecological fallacy in national culture research.” *Organization Studies* 35. 7 (2014): 1063–1086.

Cliff, Norman. “Dominance statistics: ordinal analyses to answer ordinal questions.” *Psychological Bulletin* 114. 3 (1993): 494–509.

Hartwig, Frederick, and Brian Dearing. *Exploratory Data Analysis*. Newbury Park, CA: Sage, 1979.

Heftberger, Adelheid. *Digital Humanities and Film Studies: Visualising Dziga Vertov’s Work*. Cham: Springer International Publishing, 2018.

Hodges J. L., and E.L. Lehmann. “Estimates of location based on rank tests.” *The Annals of Mathematical Statistics* 34. 2 (1963): 598–611.

Jargowsky, Paul A. “Disastrous inferences? The ecological fallacy in disaster and emergency management research.” *Disaster and Emergency Management Methods*, Collection Ed. Jason D. Rivera. New York, NY: Routledge, 2021. 48-66.

Maimets-Volt, Karie. “Arvo Pärt’s Tintinnabuli music in film.” *Music and the Moving Image* 6. 1 (2013): 55–71.

Salt, Barry. “Statistical style analysis of motion pictures.” *Film Quarterly* 28. 1 (1974): 13–22.

Salt, Barry. *Film Style and Technology: History and Analysis*, 2nd ed. London: Starword, 1992

Salt, Barry. “*Une femme mariée*.” (2007a). URL http://www.cinemetrics.lv/movie.php?movie_ID=616, accessed 1 April 2023.

Salt, Barry. “*Week-end*.” (2007b). URL http://www.cinemetrics.lv/movie.php?movie_ID=632, accessed 1 April 2023.

Sedgwick, Philip. “Understanding the ecological fallacy.” *BMJ* 351 (2015): h4773.

Vargha, András, and Harold D. Delaney. “A critique and improvement of the CL Common Language Effect Size statistics of McGraw and Wong.” *Journal of Educational and Behavioral Statistics*, 25. 2 (2000): 101–132.

Wenkert, Robert. “Reply to Emanuel Schegloff and Carlos Kruytbosch, Some comments on ‘working class authoritarianism’.” *Berkeley Journal of Sociology* 6 (1961): 109–112.

Wild, Chris. “The concept of distribution.” *Statistics Education Research Journal* 5. 2 (2006): 10–26.

Wu, Jianguo. “Scale and scaling: a cross-disciplinary perspective.” *Key Topics in Landscape Ecology*, Collection Eds. Jianguo Wu and Richard J. Hobbs. Cambridge University Press, Cambridge, 2007. 115–142.