Are Minority Opinions Shared Less?

A Conceptual Replication Using Web-Based Reviews

Birka Zapf^{1,2}, Mandy Hütter², and Kai Sassenberg^{1,2}

¹Leibniz-Institut für Wissensmedien, Tübingen, Germany

²Department of Psychology, Eberhard Karls Universität Tübingen, Germany

Abstract: Product evaluation portals on the web that collect product ratings provide an excellent opportunity to observe opinion sharing in a natural setting. Evidence across different paradigms shows that minority opinions are shared less than majority opinions. This article reports a study testing whether this effect holds on product evaluation portals. We tracked the ratings of $N = 76$ products at 12 measurement points. We predicted that the higher (lower) the mean initial rating of a product, the more positive (negative) the newly contributed ratings will differ from this baseline – as an indication of the preferred sharing of majority compared to minority opinions. We found, however, that newly added ratings were on average less extreme than earlier ratings. These results can either be interpreted as regression to the mean or evidence for the preferred sharing of minority opinions.

Keywords: online reviews, minority opinion, confirmation bias, web-based data

Product evaluations on the web provide an excellent opportunity to observe the sharing of opinions in a natural setting. Given that online shopping limits the opportunities for a direct experience with an item before buying it, consumers often rely on others' ratings and reviews (Lange, 2020). Most platforms collect ratings and reviews and provide summary statistics of past ratings. According to laboratory and interview research, opinions that disagree with the dominant opinion are less likely to be shared than those in line with it (e.g., Asch, 1956; Noelle-Neumann, 2001). However, sharing these opinions is particularly important because they add information that makes visible the wisdom of the crowd (Surowiecki, 2004). They are the basis for social change and progress, and, most important in the current context, they influence consumer decisions.

If similar effects hold for product evaluations on the web, people who hold an opinion in line with the mean rating (majority opinion) should be more likely to contribute ratings or reviews than people who hold a minority opinion (i.e., people who lean towards the opposite pole of the evaluation dimension). Thus, the mean new rating should be more extreme than the mean initial rating. Therefore, based on data from product evaluation portals, we tested whether opinions not in line with the central tendency are shared less on product evaluation portals. Thereby, we contribute not only to the understanding of the dynamics of contributions to product evaluation portals but also present

an example of how hypotheses about opinion sharing can be tested on the web.

Minority Opinions Are Shared Less Than Majority Opinions

In line with the approach often taken in survey research, we define *majority opinion* as the half of a rating scale (i.e., pro or anti) that receives more agreement (Bassili, 2003). The more the mean voiced attitude differs from the midpoint of the scale, the larger the majority in the respective sample. Research across different paradigms found that minority opinions are less likely to be shared than majority opinions. Using a scenario approach, Noelle-Neumann (2001), for instance, found that people are less willing to state their opinion publicly the more they assume that they hold a minority opinion. Similarly, Bassili (2003) found that minority opinions are expressed more slowly than majority opinions and that this effect was larger the majority was. Likewise, research on group polarization has shown that after (compared to before) hearing others' attitudes, people state opinions that lean more toward the pole of the attitude dimension that is closer to the mean of the others' attitudes (e.g., more positive on a continuum, if the average attitude is positive; e.g., Turner et al., 1989). Finally, similar effects had been documented even earlier for categorical

judgments. In Asch's (1956) classical studies, participants did not share their solution to line length comparison tasks when they were confronted with a group of people who stated a solution that was obviously wrong (for similar effects in information sharing in groups, see Stasser and Titus, 2003).

In sum, there is pervasive evidence indicating that people are more likely to share the majority compared to minority opinions. There is also evidence that such effects occur online. People are, for instance, more likely to express opinions on Facebook when they are in line with the majority opinion (Liu et al., 2017), and evidence for group polarization has also been found under conditions of mutual anonymity in computer-mediated communication – and to a stronger extent when the initial attitude indicated a clear majority (called norm in this literature Sassenberg & Boos, 2003; Spears et al., 1990).

Given that the preferred sharing of majority opinions has been shown in offline and online contexts and product evaluation portals provide users with a mean rating for each product (i.e., a clear indication of the majority opinion) as well as with the opportunity for contributing anonymously or using a pseudonym, it seems likely that the effect also generalizes to product evaluation portals. To be more precise, people should be more likely to add ratings in line with the majority than with the minority. The preferred sharing of majority compared to minority opinions should lead to newly contributed ratings that indicate a growing majority, that is, a more positive mean of the added ratings in case of a positive mean of the initial rating, and a more negative mean of the added ratings in case of a negative mean of the initial rating displayed in a product evaluation portal at a given point in time. Importantly, we assume that this polarization can be the result of the preferential sharing of majority opinions. Specifically, if primarily those people add ratings who share the opinion of the majority, their ratings should shift the distribution toward the pole of the attitude dimension that represents the majority opinion in the initial distribution.

Given the pervasive evidence that majority opinions are shared more than minority opinions, we predict that the higher (lower) the mean initial rating of a product, the more positive (negative) the newly contributed ratings will differ from this baseline. In statistical terms, this implies that a positive difference between the mean initial rating and the mean of the added ratings should be predicted by the mean initial rating. This prediction will naturally not hold for extreme cases with initial means at the poles of the attitude dimension. However, the summarized research should hold on average across initial means along the attitude dimension.

It should be noted that this prediction is based on the assumption that the ratings on a product evaluation portal

are seen as values on a continuous scale. This is indeed suggested by the mean rating communicated prominently on these portals. However, often the distribution of ratings (i.e., the relative or absolute frequency for each scale point) is just one click away from the mean rating. One could thus argue that the single rating (e.g., 4 stars), which most users opted for, can also be considered a majority rating. Given, however, that the discrete ratings are meaningless without considering the whole range of options – a 4-star rating only gains its meaning from the fact that ratings from 1 to 5 stars are possible and that more stars are better – we assume that the displayed mean rating is perceived as the majority opinion. This approach aligns with the approach often taken in survey research (Bassili, 2003; Noelle-Neumann, 2001).

Method

We tracked 87 products (e.g., toys, household items, sports equipment) on Amazon.de twice a week on 12 occasions ($t_0 - t_{11}$). Diverse products were selected covering a wide range of mean ratings ($M = 3.59$, $SD = 1.12$, range: 1.1 to 5.0 stars) and number of initial reviews ($M = 562.17$, $SD = 907.52$, range: 2 to 5,173 ratings) at t_0 . For details about the sampling procedure, see the online supplementary materials at PsychArchives (<http://dx.doi.org/10.23668/psycharchives.5061>). During the tracking period, two products were removed from Amazon.de and nine products did not receive additional ratings. The analyses below regard the remaining 76 products. Data and code are available under the following URL: <http://dx.doi.org/10.23668/psycharchives.5065>.

Prior to data collection, we conducted a pilot study in which we developed a standardized tracking procedure. Data from each product was collected twice a week at a fixed time of the day (Monday & Friday, 8 pm) for 6 weeks (on 12 occasions). At each of these points in time ($t_0 - t_{11}$), we recorded the *mean rating* of the product, the *total number of ratings*, and the *percentage of ratings* for each of the five categories. We needed to record all of this information rather than just the number of reviews and the mean rating because according to Amazon.de, the mean rating is based on a weighted mean and hence does not allow to compute the number of new items at each point in time (for an illustration of the collected data see Figure 1). It should be noted that Amazon.de states “to calculate the overall star rating and percentage breakdown by star, we don't use a simple average. Instead, our system considers things like how recent a review is and if the reviewer bought the item on Amazon. It also analyses reviews to verify trustworthiness”.

Based on the collected data, we calculated the mean of the newly *added ratings* from t_k to t_{k+l} for the measurement

Customer reviews

Figure 1. Illustration of measures assessed in product tracking.

points $k = 0 \dots 10$ and the lags $l = 1 \dots 10$. $l = 11$ was not considered because of the low number of observations included in this analysis. To be more precise, we computed the frequency of reviews for each category at k and $k + l$ by multiplying the proportion of entries in the category with the total number of entries. The number of new reviews for each category was then computed by subtracting the number at k from the number at $k + l$. Based on the five resulting scores (1–5 stars), we computed the mean of the added ratings. See Figure 2 for an illustration of this calculation.

Given that users and Amazon.de can remove ratings, the *added rating* scores entail this noise and sometimes indicate a decrease in ratings between tracking points. The relatively high frequency (twice a week) with which we collected the data allows for correcting this to some extent. Negative values were replaced by zero, given that we had no evidence that ratings had been added. On average $M = 10$ ($SD = 28$, range: 0 to 407) ratings were added per product between two subsequent measurement points. More positive scores for the *added ratings* ($M = 3.79$, $SD = 0.95$, range: 1 to 5) indicate more positive ratings.

The main dependent variable *biased sharing* ($M = -0.03$, $SD = 0.81$, range: -3.24 to 3.05) was computed by subtracting the mean initial rating at t_k from the added ratings score between t_k and t_{k+l} with $k = 0 \dots 10$ and $l = 1 \dots 10$. To support our hypothesis that people holding a majority opinion are more likely to contribute, the mean added ratings need to be more extreme than the mean initial rating.

We tested our hypothesis by regressing the *biased sharing* at t_{k+l} scores on the initial mean rating scores at t_k . In line with the recommendation of McNeish and colleagues (2017), we conducted these regressions for clustered data (by-product) using type complex in Mplus (Muthén & Muthén, 2017). In addition, we also regressed the *added rating* scores at t_{k+l} on the initial mean ratings at t_k to test the assumption that the initial mean ratings predict the mean of the added ratings (i.e., added and initial ratings are similar, and the respective product and its evaluation has not substantially changed). Given that we conducted ten tests for each of the two dependent variables, we adjusted the critical α -error from .05 to .005.

Results

We hypothesized that the initial mean ratings correlate positively with the bias scores. Contrary to this prediction, we obtained negative correlations for all time lags with standardized regression weights ranging from -0.53 to -0.31 (see Table 1, left panel). Higher initial mean ratings predicted *lower* positive deviations from these ratings. Participants did not add more ratings in the direction of the initial mean rating but fewer. Descriptively these effects became stronger across longer time intervals (based on 95% CIs but not 99.5% CIs). Figure 3 illustrates this trend: the biased sharing scores are generally larger than zero for products with an initial rating < 3 and lower than zero for products with an initial rating > 3 (72% for lag 10), indicating the sharing of more positive opinions for initially negatively evaluated products and of more negative ratings compared to the initially positively rated products (in both cases compared to the initial mean rating). In sum, our hypothesis that majority opinions are shared more than minority opinions are not supported by these data.

Given that the results did not confirm our prediction derived from earlier research, we aimed to rule out that the current findings result from specific distributions of the initial mean rating. First and foremost, for bimodal distributions, our assumption that the *mean* rating at the beginning of the interval is perceived (and functions) as the majority opinion may not be valid, given that a high relative frequency of cases does not necessarily occupy the mean. Therefore, we reran the analyses for biased sharing, *excluding* 24 products with a bimodal distribution of ratings at t_0 . Less negative correlations between initial mean ratings and biased sharing scores would suggest that our results might be driven by specific distributions. However, if anything, the opposite was the case: the standardized correlations between initial mean ratings and biased sharing ranged from -0.36 at a lag of half a week ($l = 1$) to -0.61 for a lag of 5 weeks ($l = 10$) and was, thus, descriptively even more negative than for the full sample (for details see the online supplementary material). We obtained similar results when using the median rather than the mean as a predictor. In sum, these additional analyses suggest that results (and the non-replication of earlier findings in the current context) do not result from the specific analysis strategy.

In addition, we assumed that more positive initial mean ratings are related to more positive added ratings, given that the initial and the added ratings represent the same population. In line with this assumption, the correlation between the initial mean rating and the mean of the newly added ratings was positive for all time lags ($l = 1 \dots 10$; see Table 1, right panel). The standardized regression weights ranged between 0.57 and 0.72, showing that the higher the mean rating at the beginning of an interval, the more

	t_k			t_{k+1}			$t_k - t_{k+1}$
Number of ratings (n)	200			250			
	Percentage of rating per category	<i>*n / 100</i>	Respective number of ratings	Percentage of rating per category	<i>*n / 100</i>	Respective number of ratings	
5 stars	<i>50</i>	→	100	<i>54</i>	→	135	35
4 stars	<i>20</i>	→	40	<i>20</i>	→	50	10
3 stars	<i>10</i>	→	20	<i>10</i>	→	25	5
2 stars	<i>10</i>	→	20	<i>8</i>	→	20	0
1 star	<i>10</i>	→	20	<i>8</i>	→	20	0
Mean of added ratings							4.6

mean

Figure 2. Fictitious example for the calculation of mean of added ratings. Numbers collected from Amazon.de in italics, calculation paths in gray.

Table 1. Results from regressions of biased sharing (left panel) and mean added ratings (right panel) on initial mean rating at the beginning of interval for lags of $l = 1 \dots 10$ measurement points (0.5–5 weeks). *Bs* are standardized regression weights

Lag (N)	Biased sharing				Added ratings			
	<i>B</i> (SE)	<i>Z</i>	<i>p</i>	95% CI	<i>B</i> (SE)	<i>Z</i>	<i>p</i>	95% CI
1 (597)	−0.31 (0.05)	−6.27	< .001	[−0.398; −0.218]	0.57 (0.05)	11.22	< .001	[0.470; 0.668]
2 (625)	−0.35 (0.05)	−6.27	< .001	[−0.451; −0.243]	0.61 (0.05)	11.82	< .001	[0.512; 0.715]
3 (599)	−0.38 (0.06)	−6.07	< .001	[−0.500; −0.268]	0.63 (0.06)	11.39	< .001	[0.523; 0.740]
4 (553)	−0.42 (0.07)	−5.78	< .001	[−0.547; −0.287]	0.67 (0.06)	11.24	< .001	[0.549; 0.780]
5 (497)	−0.43 (0.07)	−5.44	< .001	[−0.575; −0.287]	0.67 (0.06)	10.81	< .001	[0.551; 0.795]
6 (432)	−0.45 (0.08)	−5.07	< .001	[−0.607; −0.285]	0.68 (0.07)	10.48	< .001	[0.555; 0.810]
7 (365)	−0.46 (0.09)	−4.33	< .001	[−0.631; −0.289]	0.69 (0.07)	10.12	< .001	[0.557; 0.824]
8 (294)	−0.48 (0.09)	−3.40	< .001	[−0.661; −0.305]	0.70 (0.07)	10.39	< .001	[0.571; 0.837]
9 (222)	−0.51 (0.09)	−3.17	< .001	[−0.689; −0.336]	0.71 (0.07)	10.76	< .001	[0.581; 0.840]
10 (149)	−0.53 (0.09)	−3.00	< .001	[−0.708; −0.351]	0.72 (0.07)	11.22	< .001	[0.598; 0.851]

positive was the later contributed reviews. In other words, earlier ratings positively predict later ratings. In addition, the correlation tended to become stronger across longer time intervals (0.57 across half a week and 0.72 at lag 10, after 5 weeks). Based on 99.5% CIs, none of these differences between time intervals were statistically meaningful, based on 95% CIs, however, the correlation between initial rating and added ratings added after half a week ($l = 1$) is smaller than after 2 weeks ($l = 4$); similarly, the correlation across 1 week ($l = 2$) is smaller than the one across 5 weeks ($l = 10$). Overall, the added ratings match the initial ratings reasonably well (and more so for longer time intervals), which implies that the attitudes towards the products or the products themselves have not changed dramatically during the study.

Discussion

We aimed to replicate the established finding that minority opinions are shared less than majority opinions using web-based data from an online shopping portal. The results did not support this hypothesis. Newly added reviews did not indicate a larger majority than the initial rating, but the opposite was the case. This finding can be interpreted either as regression to the mean (i.e., *random* extreme mean ratings are followed by *valid* less extreme added ratings) or as an intentional sharing of minority opinions. It adds to existing work demonstrating general negativity biases in sharing online reviews (Godes & Silva, 2012) and towards adding less positive ratings for popular products (Le Mens et al., 2018). Minority opinions are more

Figure 3. Bias scores from ratings added within 5 weeks (lag = 10) by initial rating ($N = 149$). The diagonal lines above and below the data points mark the highest and lowest possible value, respectively.

useful as they add more information than reiterations of the majority opinion (Surowiecki, 2004). This might motivate users of product evaluation portals to share opinions deviating from the mean rating in particular. Furthermore, the authors of many ratings are given online are not identifiable. Given previous research suggests that low (own) identifiability reduces conformity to majority influence (e.g., Wu & Atkin, 2018), whereas the anonymity of others increases majority influence (for a detailed discussion, see Sassenberg & Jonas, 2007), low own identifiability in an online context might contribute to the current finding. Alternatively, the current findings might differ from research on sharing majority and minority opinions in that it often dealt with categorical opinions (e.g., decisions between discrete alternatives or opinions favoring one alternative over another in the political realm). In our analysis, we treated the product ratings as a continuous dimension and analyzed the data accordingly. However, given that survey research and the group polarization literature also dealt with attitude dimensions and found more sharing of majority than minority opinions, we do not believe this is the crucial factor. But ultimately, this and all other interpretations are highly speculative and should thus be tested in future research. In addition, given that the current study relied on a limited number of products and a single online shopping portal, the generalization of the present results is open to question, and replications are highly desirable.

Interestingly, Vinson and colleagues (2019) obtained similar effects when studying individuals' rating trajectories. In their data, individuals' own ratings correlated negatively with subsequent ratings they gave, just as in our case the ratings a product received in the past correlated negatively with subsequent ratings. This validates the general idea that new ratings are likely to be less extreme than existing ratings for the same target on product evaluation portals, be it within individuals or in a collective setting.

This study illustrates how web data can conceptually replicate well-established findings from experimental or lab research on opinion sharing in a naturalistic setting. Product evaluation portals and other social media collect and publicly document opinions on a large scale. This information can be tracked manually or automatically by crawling product evaluation portals with suitable scripts to assess newly added ratings (as in the current case) or archival data (Vinson et al., 2019). Thereby, large amounts of data can be anonymously collected (in line with common ethical guidelines) that resemble people's natural behavior and, thus, allow for the externally valid testing of psychological hypotheses.

Web-based data of this type come with challenges that should be considered when interpreting the current and similar findings. Companies at times pay people to contribute positive information about them (or even negative information about others) to product evaluation portals (Luca & Zervas, 2016). This can lead to biased reviews and contribute error variance to studies like the current one – a factor that can barely be controlled for in data collection, but platforms try to address it, for instance, by only accepting ratings and reviews from verified customers.

Further, some ratings and reviews are computer-generated or removed by the platform host, contributing to error variance. By tracking ratings at a high frequency, we were able to identify the deletion of ratings (to some extent) and reduce its impact on the results. Whether or not they succeed to do so is hard to judge, but when tracking the ratings or reviews over a longer period of time, faked ratings and reviews should mostly add error variance, and the findings should become more reliable the longer they considered time period is. This stresses the importance of multiple measurements and considering various time intervals when collecting and analyzing web-based data. Shorter intervals help to detect irregularities such as deleted ratings in our case, and longer intervals lead to more stable effects.

Depending on the research topic, the time or the number of interactions with a web platform might be the crucial factor (cf. Godes & Silva, 2012).

Especially in the context of product evaluations, most ratings are positive. Although we tried to find negatively rated products (less than two stars), many of the tracked products were still rated positively (J-shaped distribution; Hu et al., 2009). Future research should put an even stronger emphasis on the even or normal distribution of the products regarding the initial rating.

Conclusion

To conclude, web-based data provide a good opportunity to replicate well-known effects from social influence research in a naturalistic setting. In the present research, we did not find evidence for the prediction that minority opinions are shared less frequently than majority opinions, although this effect has been demonstrated across various paradigms. Rather, online raters tend to add ratings that communicate a diverging evaluation.

References

- Asch, S. E. (1956). Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychological Monographs: General and Applied*, 70(9), 1–70. <https://doi.org/10.1037/h0093718>
- Bassili, J. N. (2003). The minority slowness effect: Subtle inhibitions in the expression of Views not shared by others. *Journal of Personality and Social Psychology*, 84(2), 261–276. <https://doi.org/10.1037/0022-3514.84.2.261>
- Godes, D., & Silva, J. C. (2012). Sequential and temporal dynamics of online opinion. *Marketing Science*, 31(3), 448–473. <https://doi.org/10.1287/mksc.1110.0653>
- Hu, N., Zhang, J., & Pavlou, P. A. (2009). Overcoming the J-shaped distribution of product reviews. *Communications of the ACM*, 52(10), 144–147. <https://doi.org/10.1145/1562764.1562800>
- Lange, B. (2020, January 13). *Jeder Zweite liest Online-Bewertungen vor dem Kauf* [Every second individual reads online reviews before shopping]. Bitkom Research. <https://www.bitkom-research.de/de/pressemitteilung/jeder-zweite-liest-online-bewertungen-vor-dem-kauf>
- Le Mens, G., Kovács, B., Avrahami, J., & Kareev, Y. (2018). How endogenous crowd formation undermines the wisdom of the crowd in online ratings. *Psychological Science*, 29(9), 1475–1490. <https://doi.org/10.1177/0956797618775080>
- Liu, Y., Rui, J. R., & Cui, X. (2017). Are people willing to share their political opinions on Facebook? Exploring roles of self-presentational concern in spiral of silence. *Computers in Human Behavior*, 76, 294–302. <https://doi.org/10.1016/j.chb.2017.07.029>
- Luca, M., & Zervas, G. (2016). Fake it till you make it: Reputation, competition, and yelp review fraud. *Management Science*, 62(12), 3412–3427. <https://doi.org/10.1287/mnsc.2015.2304>
- McNeish, D., Stapleton, L. M., & Silverman, R. D. (2017). On the unnecessary ubiquity of hierarchical linear modeling. *Psychological Methods*, 22(1), 114–140. <https://doi.org/10.1037/met0000078>
- Muthén, L. K., & Muthén, B. (2017). *Mplus: Statistical analysis with latent variables user's guide* (Version 8) [Computer software]. Muthén & Muthén. http://www.statmodel.com/download/usersguide/MplusUserGuideVer_8.pdf
- Noelle-Neumann, E. (2001). *Die Schweigespirale: Theorie der öffentlichen Meinung* [Spiral of silence: Theory of public opinion] (6th ed.). Langen Müller.
- Sassenberg, K., & Boos, M. (2003). Attitude change in computer-mediated communication: Effects of anonymity and category norms. *Group Processes & Intergroup Relations*, 6(4), 405–422. <https://doi.org/10.1177/13684302030064006>
- Sassenberg, K., & Jonas, K. J. (2007). Attitude change and social influence on the net. In A. N. Joinson, K. Y. A. McKenna, T. Postmes, & U.-D. Reips (Eds.), *The Oxford handbook of Internet psychology*. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199561803.013.0018>
- Spears, R., Lea, M., & Lee, S. (1990). De-individuation and group polarization in computer-mediated communication. *British Journal of Social Psychology*, 29(2), 121–134. <https://doi.org/10.1111/j.2044-8309.1990.tb00893.x>
- Stasser, G., & Titus, W. (2003). Hidden profiles: A brief history. *Psychological Inquiry*, 14(3–4), 304–313. <https://doi.org/10.1080/1047840x.2003.9682897>
- Surowiecki, J. (2004). *The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business, economies, societies, and nations*. Random House.
- Turner, J. C., Wetherell, M. S., & Hogg, M. A. (1989). Referent informational influence and group polarization. *British Journal of Social Psychology*, 28(2), 135–147. <https://doi.org/10.1111/j.2044-8309.1989.tb00855.x>
- Vinson, D. W., Dale, R., & Jones, M. N. (2019). Decision contamination in the wild: Sequential dependencies in online review ratings. *Behavior Research Methods*, 51(4), 1477–1484. <https://doi.org/10.3758/s13428-018-1175-8>
- Wu, T. Y., & Atkin, D. J. (2018). To comment or not to comment: Examining the influences of anonymity and social support on one's willingness to express in online news discussions. *New Media and Society*, 20(12), 4512–4532. <https://doi.org/10.1177/1461444818776629>

History

Received January 28, 2021

Revision received August 9, 2021

Accepted August 9, 2021

Published online December 17, 2021

Acknowledgments

We thank Ronja Brandhorst, Cindy Hong, Claudio Pix, and Anna Vollweiter for their assistance with data collection.

Open Data

The online supplementary materials including the data and the code as part of our research are available at PsychArchives via the following URLs: <http://dx.doi.org/10.23668/psycharchives.5061> and <http://dx.doi.org/10.23668/psycharchives.5065>

Funding

Open access publication enabled by the IWM – Leibniz-Institut für Wissensmedien Tübingen.

Kai Sassenberg

Leibniz-Institut für Wissensmedien
Schleichstraße 6
72076 Tübingen
Germany
k.sassenberg@iwm-tuebingen.de