Did Einstein really say that? Testing content versus context in the cultural selection of quotations

Alberto Acerbi[1], Jamshid J. Tehrani[2]

[1] School of Innovation Science, Eindhoven University of Technology, 5600MB, Eindhoven, The Netherlands

[2] Centre for the Coevolution of Biology and Culture, Department of Anthropology, Durham University, DH1 3LE, Durham, U.K.

Did Einstein really say that? Testing content versus context in the cultural selection of quotations

Abstract

We experimentally investigated the influence of context-based biases, such as prestige and popularity, on the preferences for quotations. Participants were presented with random quotes associated to famous or unknown authors (experiment one), or with random quotes presented as popular, i.e. chosen by many previous participants, or unpopular (experiment two). To exclude effects related to the content of the quotations, all participants were subsequently presented with the same quotations, again associated to famous and unknown authors (experiment three), or presented as popular or unpopular (experiment four). Overall, our results showed that context-based biases had no (in case of prestige and conformity), or limited (in case of popularity), effect in determining participants' choices. Quotations preferred for their content were preferred in general, despite the contextual cues to which they were associated. We conclude discussing how our results fit with the well-known phenomenon of the spread and success (especially digital) of misattributed quotations, and we draw some more general implications for cultural evolution research.

Keywords

Cultural evolution, cultural transmission, context-based biases, content-based biases, quotations.

Introduction

Humans depend on social learning to acquire information and behaviours that would be otherwise difficult for individuals to learn by themselves. Theoretical models have shown that to be effective, however, social learning needs to be selective (Laland, 2004). How do we choose which ideas, beliefs and practices to adopt among the myriad of options that are available?

Research in cultural evolution suggests we use an inventory of simple heuristics, often referred to as "social learning strategies" or "cultural transmission biases", to assist our decision in respect to what, when, or from whom to copy (Boyd & Richerson, 1985; Mesoudi, 2011b). An important distinction in this inventory is made between "context-based biases" and "content-based biases" (Boyd & Richerson, 1985).

This distinction is critical because context-based biases are independent from the actual properties of the ideas or practices, whereas content-based biases, as the label suggests, refer to intrinsic characteristics of the cultural traits themselves. Examples of context-based biases are "copy prestigious individuals" (Henrich & Gil-White, 2001), "copy the majority" (Henrich & Boyd, 1998) or "copy when uncertain"(Wood et al., 2016). In all cases there is no need for the individual to directly evaluate the features of the trait to copy. If the majority is doing A in place of B, then one should copy A, no matter what A is.

Examples of content-based biases are instead "copy traits that carry survival information" (Stubbersfield, Tehrani, & Flynn, 2015) or "copy traits that elicit emotional reactions – amusement, for example" (Stubbersfield, Tehrani, & Flynn,

2017). Here the features of traits matter. Is A carrying more survival information than B?

A growing corpus of experimental studies in cultural evolution broadly supports the sketch presented above. In the case of context biases, convincing indications of, for example, the preferential copying of individual that are considered prestigious (prestige bias), have been found in laboratory (Atkisson, O'Brien, & Mesoudi, 2012; Chudek, Heller, Birch, & Henrich, 2012) as well as in ethnographic settings (Henrich & Broesch, 2011). Other experiments showed that a similar heuristic ("copy successful individuals") was used by participants to decide from whom to copy from (Mesoudi, 2011a). The empirical evidence for conformity is more scattered, but a disproportionate tendency to copy the majority (i.e. copying with a probability higher than the proportion of the majority itself, as conformity is defined in cultural evolution theory) has been found in experimental settings as well (Efferson, Lalive, Richerson, McElreath, & Label, 2008; Morgan, Rendell, Ehn, Hoppitt, & Laland, 2012; Morgan, Laland, & Harris, 2015).

Content biases have also been studied, mainly using the transmission chain (or "serial reproduction") technique (Bartlett, 1932). In these experiments, a short piece of narrative is iteratively transmitted from one participant to another. It has been found that some types of content are better remembered and repeated than others, conferring them a selective cultural advantage. In addition to the previously mentioned biases for survival-relevant information and emotional content, other content-based biases that have been studied in cultural evolution are, for example, a bias for social information (Mesoudi, Whiten, & Dunbar, 2006), a bias for minimally counterintuitive concepts – i.e. concepts that fit our intuitive cognitive expectations but with few exceptions, such as superheroes, gods, etc. (Barrett & Nyhof, 2001), or a

bias for negatively marked information (Bebbington, MacLeod, Ellison, & Fay, 2016).

One important, but hitherto largely unexplored, question concerns the relative importance of context versus content biases. What if the majority prefers A, but B carries, say, more social information than A? In what follows, we present an experiment that addresses this question. We used a sample of relatively famous quotes (such as, for example, "It is better to have loved and lost, than never to have loved at all"), and we presented them to participants. We randomly associated some of the quotes to famous authors, and others to unknown names. In a subsequent experiment we associated, again randomly, some of the quotes to a previous majority of people that preferred them, and others to a minority. We checked whether participants were inclined to prefer quotes associated to famous people, and quotes they were told were preferred by the majority.

Quotes are a useful test case, as they are relatively discrete units of cultural information that can be promptly evaluated for their content by participants, and, in the same time, are easily associated with contextual features. Context is important because quotes are usually credited to famous people, and they are commonly misattributed. The quote "The definition of insanity is doing the same thing over and over again and expecting a different result", for example, is often incorrectly attributed to Albert Einstein. However, the earliest exact match of the quote appears in a Narcotics Anonymous information pamphlet in 1981, some 25 years after Einstein's death[1]. The fact that most people attribute the quote to Einstein rather than its true source is suggestive of the value added by fame to the "quotability" of a phrase. On the other side, content is important because there must be something

[1] http://quoteinvestigator.com/2017/03/23/same/

particularly appealing about the specific message in a quote – we don't just quote anything and everything that a famous person has said. A recent study by Lerique and Roth (2017), for example, provides intriguing evidence for content-biased transmission in quotations, showing that quotes copied from one website to another tended to be transformed according to predictable rules, for example replacing difficult words with simpler synonyms.

In the experiments, we address the following questions:

1) Does being associated to a famous author influence whether a quote is more liked? In addition: does the domain of the quote modulate this influence? Our hypothesis was that the influence of the association of a prestigious author with a quote would have been stronger when the author was known as an "expert" of the quote's domain (hence the topic "Science" and "Literature" associated with famous scientists and writers), less strong when the domain of the quote was "Money" or "Success" (for which *any* famous author could know more than the average people, without being experts in the domain), and finally even less for domains, such as "Love" and "Friendship", that could be intuitively considered common knowledge.

2) Does the popularity of a quote influence whether people like the quote? We tested here two different hypotheses. The first one is that people would be conformist in the technical sense defined above, i.e. that they would disproportionally (with a probability higher than the popularity of the quote itself) prefer a popular quote. The second – weaker – hypothesis is that popular quote would simply be more preferred than unpopular one. In addition: does the domain of the quote modulate this influence? We reasoned that people might attend more to popularity in domains that do not require expert knowledge, such as "Love" and "Friendship" than ones like "Science" and "Literature", or "Money" and "Success", where common knowledge

might be an unreliable guide to the usefulness of the information contained in the quote.

General methods

We carried out four experiments. In the first two experiments (experiment one and two), randomly extracted pairs of quotes of the same domain were assigned to participants. In experiment one famous and unknown author were assigned to the quotes, while in experiment two one quote was presented as "popular" and one was not. Participants were asked to choose the quote they preferred in the pair.

In the other two experiments (experiment three and four) all participants were presented with the same quotes. The quotes were associated alternatively with famous or unknown authors (experiment three) or were presented as popular or unpopular (experiment four). Participants were asked to rate how much they liked each quote.

Selection of quotes

We selected from the website http://www.quotationspage.com 10 quotes for each of these six topics: "love", "friendship", "money", "success", "science", and "literature". We chose quotes that were, according to our judgment, not particularly recognizable, so that assigning to them an unknown – or wrong – author would not jar with participants' prior knowledge about sources. We also chose 4 quotes to use as a "distractor", and two quotes to use as a "control" (see below). All quotes were a single sentence statement, to avoid any bias related to length. The list of the 66 quotes used in the experiment is provided in Supplementary Material (quotations.pdf).

Content only evaluation

We recruited 200 participants through crowdflower.com. Each participant was paid 1.00$ to carry out the task, which took between two and three minutes to complete (143 seconds on average). After completing the task, the participants were debriefed about the aims of the experiment and given the option to withdraw their data. None of the participants chose this option. Participants were also informed that some quotes in the experiment were misattributed and provided a link to the website where the quotes (and authentic sources) were sourced from. We followed this procedure for all the experiments described below. The University of Bristol granted the ethical approval for the experiment.

Participants were asked to help us to "Choose the most inspirational quote" and presented a questionnaire with seven questions. Each of the seven questions included a pair of quotes, and the participant was asked to choose the one s/he preferred between the two - see screenshot in Supplementary Material (screen1.pdf). Six questions concerned the six topics above. Each question presented two quotes randomly selected in each topic. One "Control" question always presented the same two quotes (randomly associated to a famous and to an unknown author). One of the quotes was meaningless ("The it then said it to the boring good morning"), and the participants preferring this quote were excluded from the analysis. Finally, the order of presentation of the quotes was randomized for each participant.

We collected data from 174 valid participants (26 being excluded because of the wrong answer in the "control" question). Each of the 60 quotes was presented on average 34.8 times (SD=5.0, max=46, min=24).

Selection of "famous" authors

We first extracted names to use as "famous" authors from the Pantheon 1.0 dataset (Yu, Ronen, Hu, Lu, & Hidalgo, 2016; available online at: pantheon.media.mit.edu). We considered names from the category "people", with any place of birth, and born between 1800 and 2010. We extracted 20 names from the Pantheon domain "All" (including personalities from all domains), 10 from the domain "Humanities", and 10 from the domain "Science & Technology". We excluded, in order to avoid biases, women (only Marie Curie, in the "Science & Technology" domain, and Marilyn Monroe, in the "All" domain, were found) and possibly controversial political figures from the "all" domain (Karl Marx, Adolf Hitler, Che Guevara, Joseph Stalin, Vladimir Lenin, Benito Mussolini, and Mao Zedong). We obtained a total of 30 different famous names (as 10 famous authors were repeated in different domains, for example Albert Einstein was present both in the "All" and in the "Science & Technology" domain).

We tested if famous names were indeed recognised as such by participants, contrasting them with a sample of 30 randomly generated male names ("Unknown" sample) that was then used for the experiments. The list of famous and unknown names is provided in Supplementary Material (authors.pdf). Data were collected from 100 participants recruited through crowdflower.com. Each participant was paid 0.40$ to complete the task. The task took approximate one minute to complete (62 seconds on average). Participants were asked to help us to "Rate how famous (well-known) contemporary or past celebrities are". Each participant was presented with four names, chosen at random in each category ("All", "Science and Technology", "Humanities", "Unknown"). Each name was presented with a multiple-choice question ("How famous do you think he is?" with possible answers: "very famous", "famous", "a little famous", "not famous at all").

There was a significant difference in the rating of unknown and famous names in all three categories, demonstrating that participants recognized as famous the names we extracted from the Pantheon dataset, and not the random names. Small variations were present in different domains (for example, the names from the "Science & Technology" domain were known slightly better than the names from the "Humanities" domain), but all the differences with the unknown names were significant at the same level (Wilcoxon signed-rank test, all $p<0.0001$, all $N=100$, see also Supplementary Material – additional_stats_info.pdf, Table 1 SM),

Experiment one: famous versus unknown authors

Methods

We recruited 200 participants through crowdflower.com. Each participant was paid 1.30$ to complete the task (average task duration: 152 seconds). As above, subjects were asked to help us to "Choose the most inspirational quote" and presented a questionnaire with nine questions. Each of the nine questions included a pair of quotes, and the participant was asked to choose the one s/he preferred between the two. Table 1 shows how quotes and authors were assigned to participants. For each topic, two random quotes were selected, and authors from the samples described in Table 1 were also randomly extracted.

Two questions – not used in the analysis – included two random quotes both associated to famous or unknown names, respectively. The rationale for including these two "Distractors" was to avoid participants realizing the hypothesis that we were testing (which may have been obvious if all the questions pitted a quote by one famous and one by an unknown author). Finally, the order of presentation of the

quotes, as well as the order of the authors inside each questions, was randomized for each participant.

Table 1 here

Results

We excluded 39 participants due to preferring the meaningless quote in the control question, remaining with 161 valid participants. We calculated, for each of the possible 60 quotes, how many times a quote was preferred when associated to a famous author (hence "opposing" a quote of the same topic, associated to an unknown author), and how many times it was preferred when associated to an unknown author. Each quote was presented on average 32.2 times overall to the 161 valid participants (SD=5.1, max=48, min=24). We performed, for the three separate categories of topics ("Love/Friendship", "Money/Success", "Science/Literature"), three separate Wilcoxon signed-ranked tests, comparing the success rate of quotes associated to famous and to unknown authors. All tests gave non-significant results (Wilcoxon signed-rank test, p=0.11, p=0.42, p=0.20, all N=20, see also Supplementary Material – additional_stats_info.pdf, Table 2 SM), indicating that participants did not preferred a quote when it was associated with a famous authors more than when it was associated to an unknown author (see Figure 1).

Figure 1 here

To test the effect of the content, we used the results provided in the "Content only Evaluation" (see above) as one of the predictors of two linear models. The first

linear model included, as a response, the success rate of quotes in experiment one (famous versus unknown authors), and, as the other predictor, the proportion of times the quote was associated to a famous author in experiment one. The model was overall significant ($p<0.001$, $R^2=0.36$), and showed that the proportion of wins in "Content only Evaluation" ($p<0.001$, $t=5.69$), but not the proportion of times the quote was associated to a famous author in experiment one ($p=0.66$, $t=-0.44$) explained the success in experiment one (see also Figure 2). In other words, participants evaluated the content of the quotes, but not the fact that they were associated to a famous author, to choose among them in the "famous versus unknown authors" experiment.

Figure 2 here

Experiment two: popular versus unpopular quotes

Methods

The structure of experiment two was analogous to experiment one, but instead of authors, quotes were associated with a popularity score ("N people already chose this quote"). Using the same arrangement of Table 1, the quote associated to a "Famous" author was now a "Popular" quote, while the quote associated to an "Unknown" author was, in experiment two, an "Unpopular" quote. All quotes and their order were randomized again for experiment two. The numbers of people that already chose "Popular" and "Unpopular" quotes were randomly generated with the constraint that, for each question, unpopular quotes were assigned a random number of people that already chose them between 100 and 1000, and popular quotes were

presented as chosen approximately by three times more people than unpopular ones. Following the logic of experiment one, the two quotes in the Distractors were presented as chosen approximately by the same number of people. As in experiment one, we recruited 200 participants through crowdflower.com. Each participant was paid 1.30$ to complete the task (average task duration: 146 seconds).

Results

We analysed the answers from 165 participants (35 were excluded). Each of the 60 quotes was presented on average 33.0 times (SD=4.6, max=46, min=21). We first checked if participants showed any conformist tendency. A visual inspection of the data (see Figure 3) clearly shows that this was not the case. To show a *disproportionate* tendency to prefer popular quotes, participants should have preferred them with a probability higher than the frequency they were presented (3/4 of the total presumed preferences, see Methods above). Similarly, unpopular quotes should have been preferred with a probability lower than the frequency presented. In Figure 3, the shaded area of the plots represents these hypothetical outcomes.

Figure 3 here

In the subsequent analysis we focused on whether popularity still had some effect on participants' evaluations, even if it was not "conformist" in the technical sense examined above. The same analysis of experiment one was conducted for experiment two (popular versus unpopular quotes). Three Wilcoxon signed-ranked tests gave here a significant difference between the proportions of wins of “popular” versus “unpopular” quotes (Wilcoxon signed-rank test, all $p < 0.001$, all $N=20$, see

also Supplementary Material – additional_stats_info.pdf, Table 3 SM), indicating that participants preferred "popular" quotes (see Figure 4). As we did not have specific hypotheses on the role of topic domains for popularity, we did not analyse possible differences in the results between the three categories of topics.

Figure 4 here

According to the same logic applied to experiment one, we ran a linear model, in which the response variable was the success rate of quotes in experiment two (popular versus unpopular quotes), and two predictors were used: the proportion of wins in the "Content only evaluation" test and the proportion of times the quote was "popular" in experiment two. The model was again overall significant ($p<0.001$, $R^2=0.47$), but, differently from experiment one, showed that both the proportion of wins in experiment the "Content only evaluation" test ($p<0.001$, $t=6.28$), and, to a lesser degree, the proportion of times the quote was popular in experiment two ($p<0.005$, $t=3.37$) explained the success in experiment two (see also Figure 5).

Figure 5 here

Experiment three: Single quotes and fame

Methods

To avoid any effect of the content of quotes, we ran a second series of experiments, in which all participants were presented with the same quotes, and the only variation was the fact that they were associated with Famous or Unknown

authors (Experiment three) or were considered Popular or Unpopular (Experiment four)

For experiment three we recruited 200 participants through crowdflower.com. Each participant was paid 0.70$ to complete the task (average task duration: 148 seconds). Participants were again asked to help us to "Choose the most inspirational quote", presented a questionnaire with seven quotes, and informed of the experiment after completing the task, as described above. Each quote was presented with a multiple-choice question ("How good do you think this quote is?" with possible answers: "very good", "good", "average", "not particularly good").

All participants were assigned the same seven quotes, six for each of the possible topics, plus the same "Control" quote described above (see all quotes in Table 2). The data of participants that answered that the meaningless control quote was "very good", "good", or "average" were discarded. For each of the quotes, half of the participants were randomly assigned a famous author (from the sample "All", or from the sample "Science and Technology" for the topic "Science", and from the sample "Humanities" for the topic "Literature", analogously to experiment 1), and the other half of participants was assigned a name from the "Unknown" sample. The order of presentation of the quotes was finally randomised.

Table 2 here

Results

We discarded 10 participants that evaluated positively the control quote, remaining with 190 valid subjects. For each topic, we compared the evaluations of the quote associated with the famous author with the evaluations of the quote associated

with the unknown name. While the former where indeed higher (see Figure 6, upper panel), the differences were not significant (Wilcoxon rank-sum test, all $p>0.05$, see also Supplementary Material – additional_stats_info.pdf, Table 4 SM), consistently with the results of experiment one.

Experiment four: Single quotes and popularity

Methods

As above, we kept the same structure of experiment three, and we replaced "Famous" and "Unknown" authors with "Popular" and "Unpopular" quotes ("N people think this is a good quote"). The number of people that already chose "Popular" and "Unpopular" quotes were generated by selecting a random number between 100 and 1000 for each participant and by multiplying this number by 0.75 for popular quotes and by 0.25 for unpopular ones (adding randomness). In this way we kept the approximate ratio 1/3 between people who chose popular and unpopular quotes present in Experiment two. As in experiment three, we recruited 200 participants through crowdflower.com. Each participant was paid 0.70$ to complete the task (average task duration: 154 seconds).

Results

We retained the answers of 198 participants, and compared the evaluations of the quote presented as "popular" versus the evaluations of the quote presented as "unpopular". As above, "popular" quotes were rated higher than the same quotes, presented as "unpopular" (see Figure 6, lower panel). The difference was significant

for two topics, “Friendship” (Wilcoxon rank-sum test, $p<0.005$) and “Science” (Wilcoxon rank-sum test, $p<0.05$) and not significant for the others (see also Supplementary Material – additional_stats_info.pdf, Table 5 SM),

Figure 6 here

Discussion

Our experiments gave some indication, contrary to our expectations, that context-based cultural transmission biases had less effect than the actual content in determining how participants evaluated the material presented. The first experiment showed that the fact that a quote was associated or not to a famous author was not important in determining whether it was preferred or not. The second experiment showed both that our participants were not conformist – in the technical sense defined in cultural evolution, i.e. having a *disproportionate* tendency to copy the majority – and that, while the perceived popularity of a quote had an effect on their choices, this effect was relatively small in respect to the effect of the content of the quote itself. Finally, experiments three and four showed that, when controlling for the content by presenting the same quote to participants, popularity and prestige had, again, a limited effect. We found two significant differences in experiment four, showing that participants preferred consistently the popular quote in the domains of “Friendship” and “Science”. However, the effect was present in only two of six domains, and we did not have theoretical reasons to expect that “Friendship” and “Science” would show a bigger influence of a popularity bias. We tentatively interpret these two

significant results only suggesting, consistently with the results of experiment two, that *some* effect of popularity was present, more than in the case of fame/prestige.

These results may seem surprising, given the apparently common tendency for people to misattribute quotes to famous people (recall our earlier example of quoting Einstein, rather than Narcotics Anonymous, on the relationship between repetition and insanity). While, at first sight, this phenomenon would seem to exemplify prestige bias, our results suggest that other explanations should be considered. For example, it could be that people remember the content better than who said it, so when they re-tell or "share" the quote they could make errors in attribution. The aggregation of these errors is likely to lead to more quotes being misattributed to famous people simply because they are better known (so Einstein is bound to pick up more accidental misattributions than, say, Bohr, simply because fewer people would know or remember who Bohr was). According to this interpretation, the success of quotations would not be the result of being misattributed to famous authors. On the contrary, misattributions would be the result of the wide diffusion of "good" quotations.

On a more general level, we may ask how the results of our experiments can contribute to the broad field of cultural evolution. There are two important features of the experiment that need to be considered to evaluate the scope of our results. First, no expertise was required to choose between the alternatives. A basic tenet of cultural evolution theory is that social information is valuable when individual information is costly and/or difficult to obtain ("costly information hypothesis" in Boyd & Richerson, 1985). This was clearly not the case in our scenario, so that it is likely this may explain why participants did not consider the social cues that were provided with the quotes (for recent experiments showing the relationship between task difficulty

and (under)use of social information see e.g. Acerbi, Tennie, & Mesoudi, 2016; Morgan et al., 2012)

The second important feature however was that the choice was, for the participants, completely cost-free. In this case, previous studies indicate that context-based biases are expected to have an important role. To limit to examples that directly refer to cultural evolution theory, Coultas (2004) found that university students were influenced by the majority (but not conformist) about seemingly irrelevant choices such as writing a date analogically ("2 February 2017") or numerically ("2/2/2017"), or covering or not the keyboard of the public computer they used. Claidière et al. (2014) showed that the visitors of a zoo, given the opportunity to answer to questions on a card in exchange of a small prize, wrote (or drew) their contribution according to what they perceived others visitors did previously. One of the illustrations of prestige-based bias used in cultural evolution, that is, the influence of stars like Michael Jordan in advertisement (Boyd, Richerson, & Henrich, 2011), involves a "task" (choosing the underwear's brand) that is cost-free and does not require expertise, hence it is quite similar to our scenario. Future studies should systematically test how the variation on the two axes of task difficulty and task importance may influence the usage of context-based transmission biases and social cues in general.

Our results contribute to a growing body of works that found contrasting results on the effects of context-based biases. For example, Salganik, Dodds, & Watts (2006) produced results very similar to our study. Salganik et al. (2006) created an "artificial market" were individuals could download previously unknown songs and, in the "social-influence" condition, see how many times the songs have been previously downloaded. While the study is often cited to support the importance of the influence of popularity on individual choices, Salganik et al. (2006) found that there was a

strong correlation between the success of songs in the "social influence" condition and in the control condition, where individuals did not have contextual cues of popularity, mirroring what happened in our experiments. Notice that, also in this case, the choice (downloading or not a song) was low-cost and did not require previous experience. Similarly, Priestley & Mesoudi (2015) studying the behaviour of users of the aggregator website Reddit.com, found that social influence (users are more likely to up-vote content that others have previously up-voted) had a smaller effect than expected.

Establishing the relative importance of context and content biases, for cultural evolutionary studies, is a task that goes beyond the mere need for terminological precision. Context-based biases are relatively simple, domain-general, heuristics. If they are the main driving force of cultural evolution, cultural evolution studies should mainly focus on population-level dynamics. Modelling strategies, or theoretical approaches, in which the cognitive properties of human individuals are only minimally sketched will do the job. On the contrary, content-based biases depend on domain-specific cognitive aspects, and, if the success of practices and ideas depend mostly on those, cultural evolutionists need to pay particular attention to the subtleties of human cognition.

Acknowledgements

References

Acerbi, A., Tennie, C., & Mesoudi, A. (2016). Social learning solves the problem of narrow-peaked search landscapes: experimental evidence in humans. *Royal Society Open Science*, *3*(9). http://doi.org/10.1098/rsos.160215

Atkisson, C., O'Brien, M. J., & Mesoudi, A. (2012). Adult learners in a novel environment use prestige-biased social learning. *Evolutionary Psychology*, *10*(3), 519–37.

Barrett, J., & Nyhof, M. (2001). Spreading Non-natural Concepts: The Role of Intuitive Conceptual Structures in Memory and Transmission of Cultural Materials. *Journal of Cognition and Culture*, *1*(1), 69–100. http://doi.org/10.1163/156853701300063589

Bartlett, F. C. (1932). *Remembering : a study in experimental and social psychology*. Cambridge: Cambridge University Press.

Bebbington, K., MacLeod, C., Ellison, T. M., & Fay, N. (2016). The sky is falling: evidence of a negativity bias in the social transmission of information. *Evolution and Human Behavior*. http://doi.org/10.1016/j.evolhumbehav.2016.07.004

Boyd, R., & Richerson, P. J. (1985). *Culture and the Evolutionary Process*. Chicago: University of Chicago Press.

Chudek, M., Heller, S., Birch, S., & Henrich, J. (2012). Prestige-biased cultural learning: bystander's differential attention to potential models influences children's learning. *Evolution and Human Behavior*, *33*(1), 46–56. http://doi.org/10.1016/j.evolhumbehav.2011.05.005

Claidière, N., Bowler, M., Brookes, S., Brown, R., & Whiten, A. (2014). Frequency of behavior witnessed and conformity in an everyday social context. *PloS*

One, *9*(6), e99874. http://doi.org/10.1371/journal.pone.0099874

Coultas, J. C. (2004). When in Rome... An Evolutionary Perspective on Conformity. *Group Processes & Intergroup Relations*, *7*(4), 317–331. http://doi.org/10.1177/1368430204046141

Efferson, C., Lalive, R., Richerson, P., McElreath, R., & Label, M. (2008). Conformists and mavericks: the empirics of frequency-dependent cultural transmission. *Evolution and Human Behavior*, *29*(1), 56–64. http://doi.org/10.1016/j.evolhumbehav.2007.08.003

Henrich, J., & Boyd, R. (1998). The Evolution of Conformist Transmission and the Emergence of Between-Group Differences. *Evolution and Human Behavior*, *19*(4), 215–241. http://doi.org/10.1016/S1090-5138(98)00018-X

Henrich, J., & Broesch, J. (2011). On the nature of cultural transmission networks: evidence from Fijian villages for adaptive learning biases. *Philosophical Transactions of the Royal Society B*, *366*(1567), 1139–48. http://doi.org/10.1098/rstb.2010.0323

Henrich, J., & Gil-White, F. J. (2001). The evolution of prestige: freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behavior*, *22*(3), 165–196. http://doi.org/10.1016/S1090-5138(00)00071-4

Laland, K. N. (2004). Social learning strategies. *Learning & Behavior*, *32*(1), 4–14. http://doi.org/10.3758/BF03196002

Lerique, S., & Roth, C. (2017). The Semantic Drift of Quotations in Blogspace: A Case Study in Short-Term Cultural Evolution. *Cognitive Science*, online first. http://doi.org/10.1111/cogs.12494

McElreath, R., Lubell, M., Richerson, P. J., Waring, T. M., Baum, W., Edsten,

E., … Paciotti, B. (2005). Applying evolutionary models to the laboratory study of social learning. *Evolution and Human Behavior*, *26*(6), 483–508. http://doi.org/10.1016/j.evolhumbehav.2005.04.003

McGraw, K.O., &Wong, S.P. (1992) A Common Language Effect Size Statistic. *Psychological Bulletin*, 111, 361-365.

Mesoudi, A. (2011a). An experimental comparison of human social learning strategies: payoff-biased social learning is adaptive but underused. *Evolution and Human Behavior*, *32*(5), 334–342. http://doi.org/10.1016/j.evolhumbehav.2010.12.001

Mesoudi, A. (2011b). *Cultural Evolution*. Chicago: University of Chicago Press.

Mesoudi, A., Whiten, A., & Dunbar, R. (2006). A bias for social information in human cultural transmission. *British Journal of Psychology*, *97*(3), 405–423. http://doi.org/10.1348/000712605X85871

Morgan, T. J. H., Laland, K. N., & Harris, P. L. (2015). The development of adaptive conformity in young children: effects of uncertainty and consensus. *Developmental Science*, *18*(4), 511–24. http://doi.org/10.1111/desc.12231

Morgan, T. J. H., Rendell, L. E., Ehn, M., Hoppitt, W., & Laland, K. N. (2012). The evolutionary basis of human social learning. *Proceedings. Biological Sciences / The Royal Society*, *279*(1729), 653–62. http://doi.org/10.1098/rspb.2011.1172

Priestley, M., & Mesoudi, A. (2015). Do Online Voting Patterns Reflect Evolved Features of Human Cognition? An Exploratory Empirical Investigation. *PloS One*, *10*(6), e0129703. http://doi.org/10.1371/journal.pone.0129703

Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, *311*(5762),

854–6. http://doi.org/10.1126/science.1121066

Stubbersfield, J. M., Tehrani, J., & Flynn, E. G. (2017). Chicken Tumours and a Fishy Revenge: Evidence for Emotional Content Bias in the Cultural Transmission of Urban Legends. *Journal of Cognition and Culture*. http://doi.org/10.13140/RG.2.1.5181.0967

Stubbersfield, J. M., Tehrani, J. J., & Flynn, E. G. (2015). Serial killers, spiders and cybersex: Social and survival information bias in the transmission of urban legends. *British Journal of Psychology (London, England : 1953)*, *106*(2), 288–307. http://doi.org/10.1111/bjop.12073

Wood, L. A., Harrison, R. A., Lucas, A. J., McGuigan, N., Burdett, E. R. R., & Whiten, A. (2016). "Model age-based" and "copy when uncertain" biases in children's social learning of a novel task. *Journal of Experimental Child Psychology*, *150*, 272–284. http://doi.org/10.1016/j.jecp.2016.06.005

Yu, A. Z., Ronen, S., Hu, K., Lu, T., & Hidalgo, C. A. (2016). Pantheon 1.0, a manually verified dataset of globally famous biographies. *Scientific Data 2*,150075. http://doi: 10.1038/sdata.2015.75

Figures captions

Figure 1: Comparison of quotes' proportion of wins across the three topic groups in experiment one (Famous versus Unknown authors). Boxplots show medians and interquartile ranges, with whiskers extending to 1.5*IQR.

Figure 2: Fame and content versus quotes' success. Left panel: Linear regression of the proportion of times a quote was associated with a famous author in experiment one versus the proportion of wins in experiment one. The shaded area shows the 95% confidence interval. Right panel: Linear regression of the proportion of wins in experiment "Content only Evaluation" versus the proportion of wins in experiment one. The shaded area shows the 95% confidence interval.

Figure 3: Average proportion of wins across the three topic groups in experiment two (Popular versus unpopular quotes) versus the frequency they were presented to subjects. Bars represented standard deviations of the data. The shades areas of the plots show where data points would have been expected, if participants had shown a conformist tendency.

Figure 4: Comparison of quotes' proportion of wins across the three topic groups in experiment two (Popular versus unpopular quotes). Boxplots show medians and interquartile ranges, with whiskers extending to 1.5*IQR.

Figure 5: Popularity and content versus quotes' success. Left panel: Linear regression of the proportion of times a quote was presented as "popular" in

experiment two versus the proportion of wins in experiment two. The shaded area shows the 95% confidence interval. Right panel: Linear regression of the proportion of wins in the "Content only evaluation" test versus the proportion of wins in experiment two. The shaded area shows the 95% confidence interval.

Figure 6: Comparison of quotes' success across the six topics in experiment three and four. Upper panel: Percentage of success calculated as CLES ("common language effect size" McGraw & Wong, 1992; i.e. how many times, given all possible pairings, the quote in one condition was evaluated higher than the same quote in the other condition) across topics in experiment three (Single quotes and fame). Notice the sum for each topic is not 100, as a proportion of pairings resulted in ties. Lower panel: Percentage of success calculated as CLES across topics in experiment four (Single quotes and popularity).

Table 1

TOPIC	AUTHOR 1	AUTHOR 2
Love	Famous “All”	“Unknown”
Friendship	Famous “All”	“Unknown”
Money	Famous “All”	“Unknown”
Success	Famous “All”	“Unknown”
Science	Famous “Science and Technology”	“Unknown”
Literature	Famous “Humanities”	“Unknown”
Distractor 1	“Unknown”	“Unknown”
Distractor 2	Famous “All”	Famous “All”
Control	Famous “All”	“Unknown”

Table 1: How quotes and authors were presented to participants in experiment one.

TOPIC	QUOTE
Love	It is better to have loved and lost, then never to have loved at all
Friendship	The meeting of two personalities is like the contact of two chemical substances: if there is any reaction, both are transformed
Money	One of the greatest disservices you can do to a man is to lend him money that he can't pay back
Success	If you can break down those walls you've spent so many years building to protect yourself, you can achieve anything
Science	Science may set limits to knowledge, but should not set limits to imagination
Literature	The man who does not read good books has no advantage over the man who can not read them
Control	The it then said it too the boring good morning

Table 2: Quotes used in experiments three and four.

Figure 1

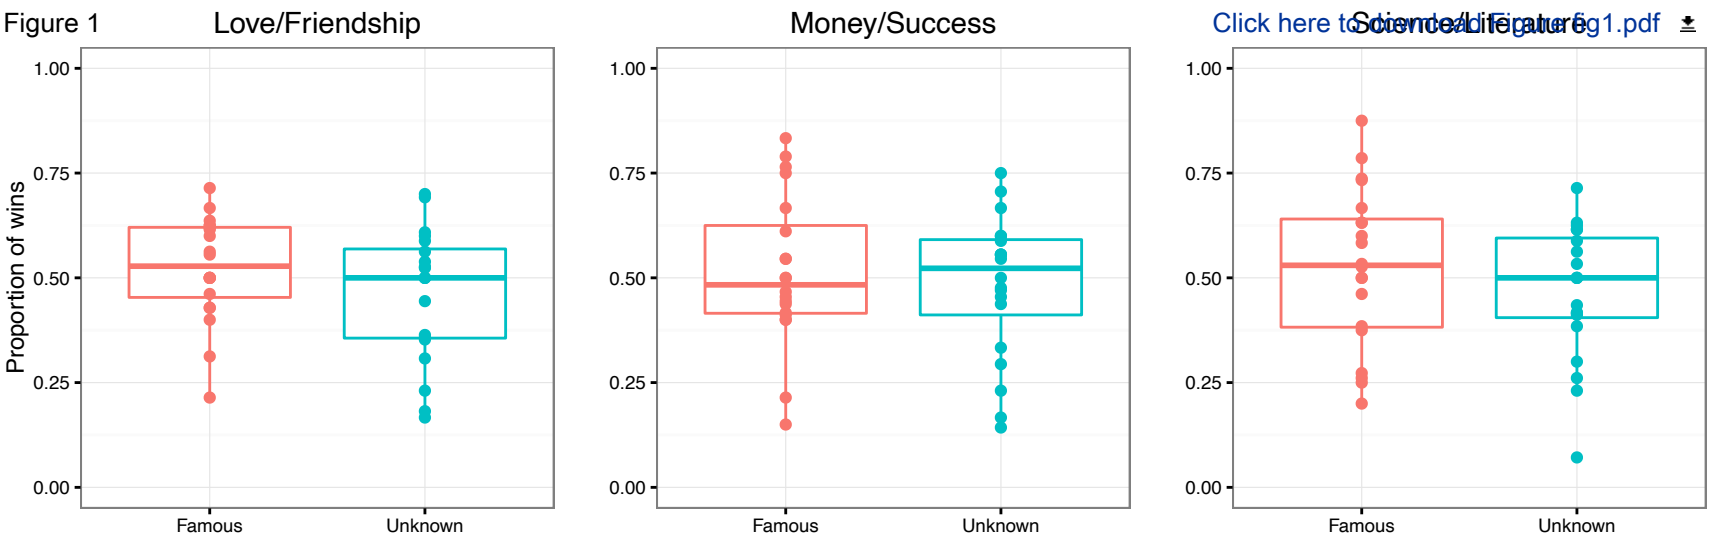

Click here to download Figure fig1.pdf

Figure 2

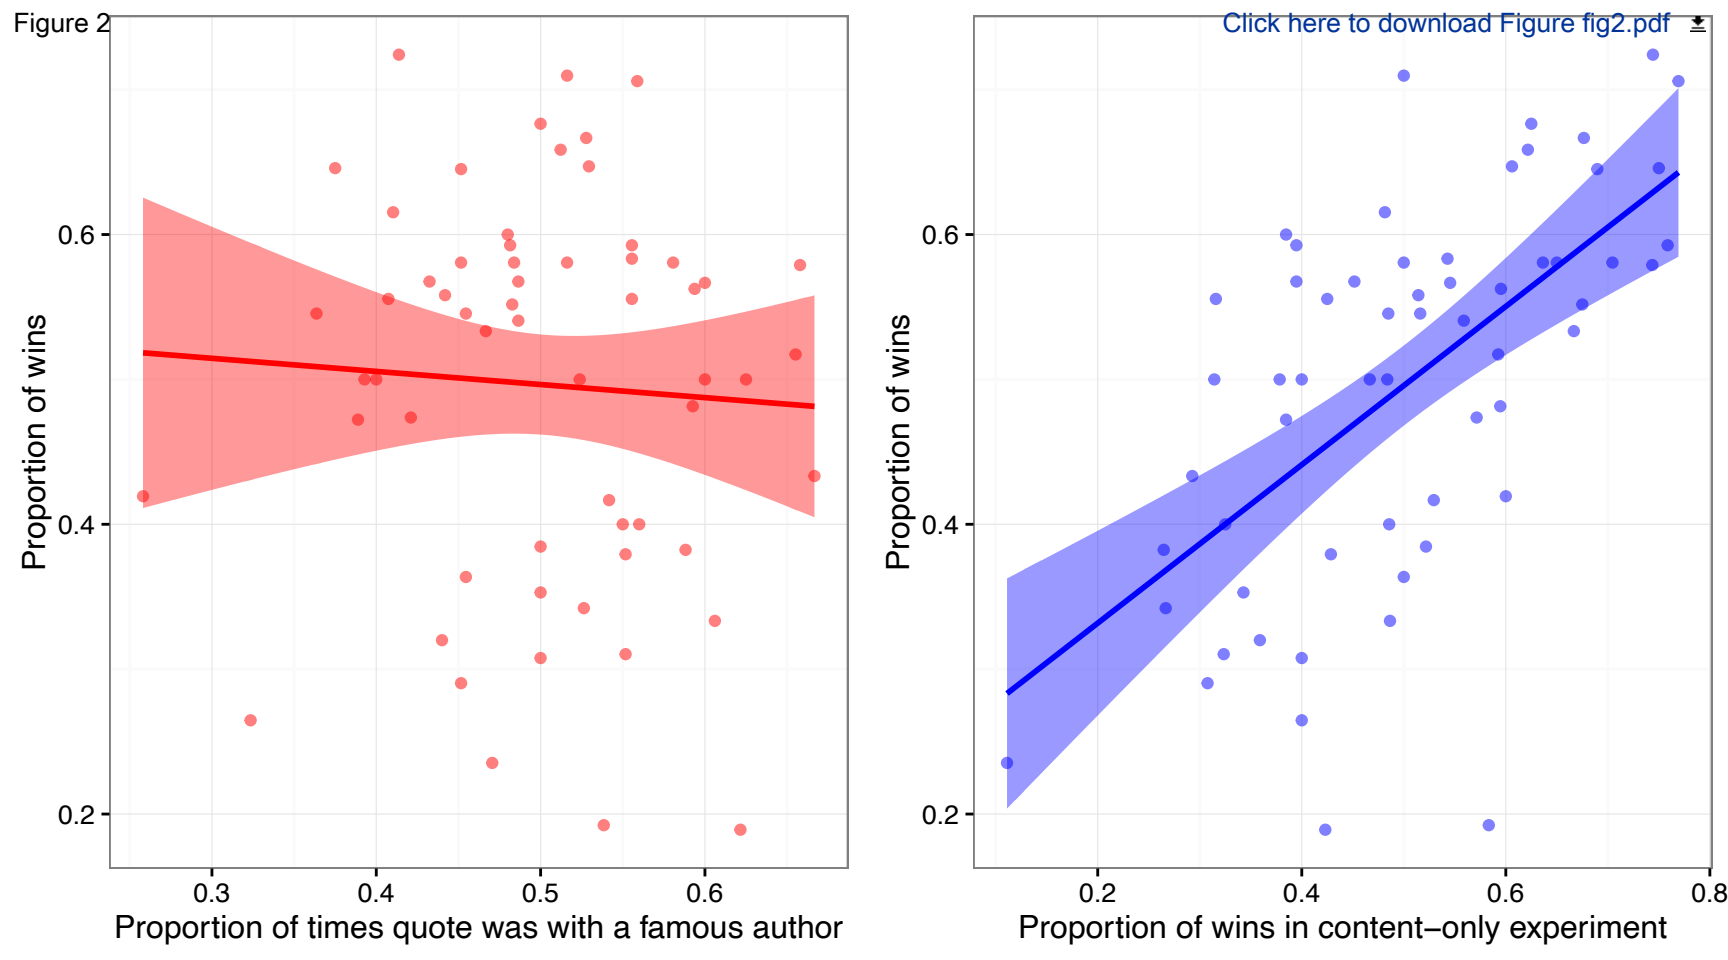

Figure 3

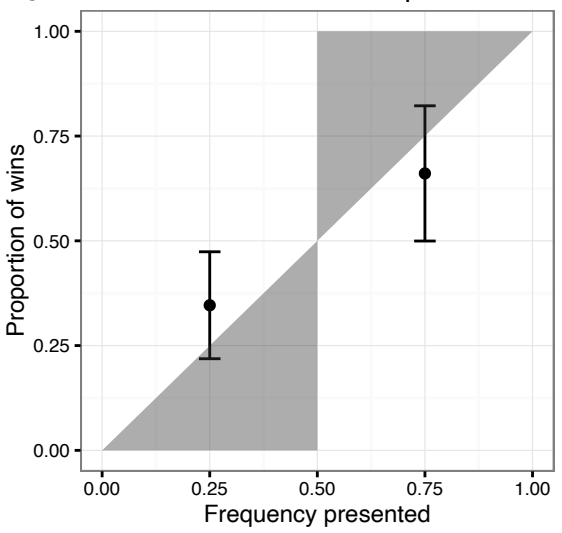
Love/Friendship
Proportion of wins
1.00
0.75
0.50
0.25
0.00
0.00
0.25
0.50
0.75
1.00
Frequency presented

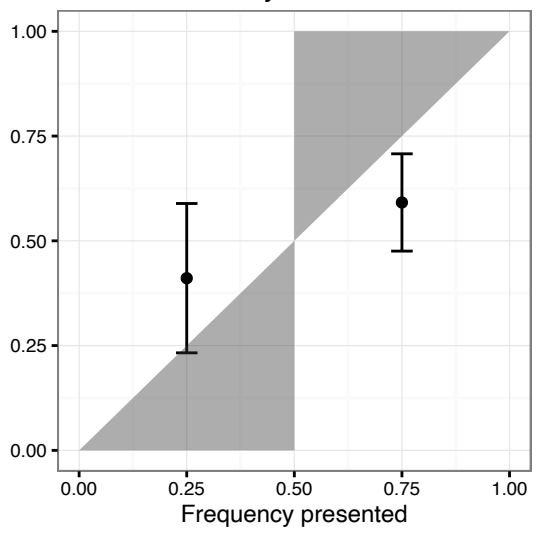
Money/Success
1.00
0.75
0.50
0.25
0.00
0.00
0.25
0.50
0.75
1.00
Frequency presented

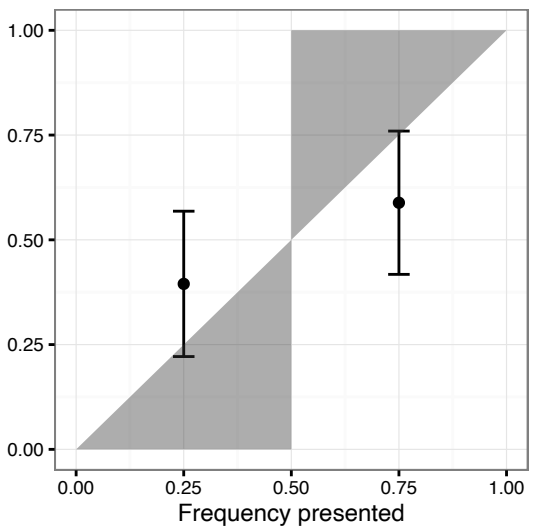
Click here to download Figure fig3.pdf
Social/Lifestyle
1.00
0.75
0.50
0.25
0.00
0.00
0.25
0.50
0.75
1.00
Frequency presented

Figure 4

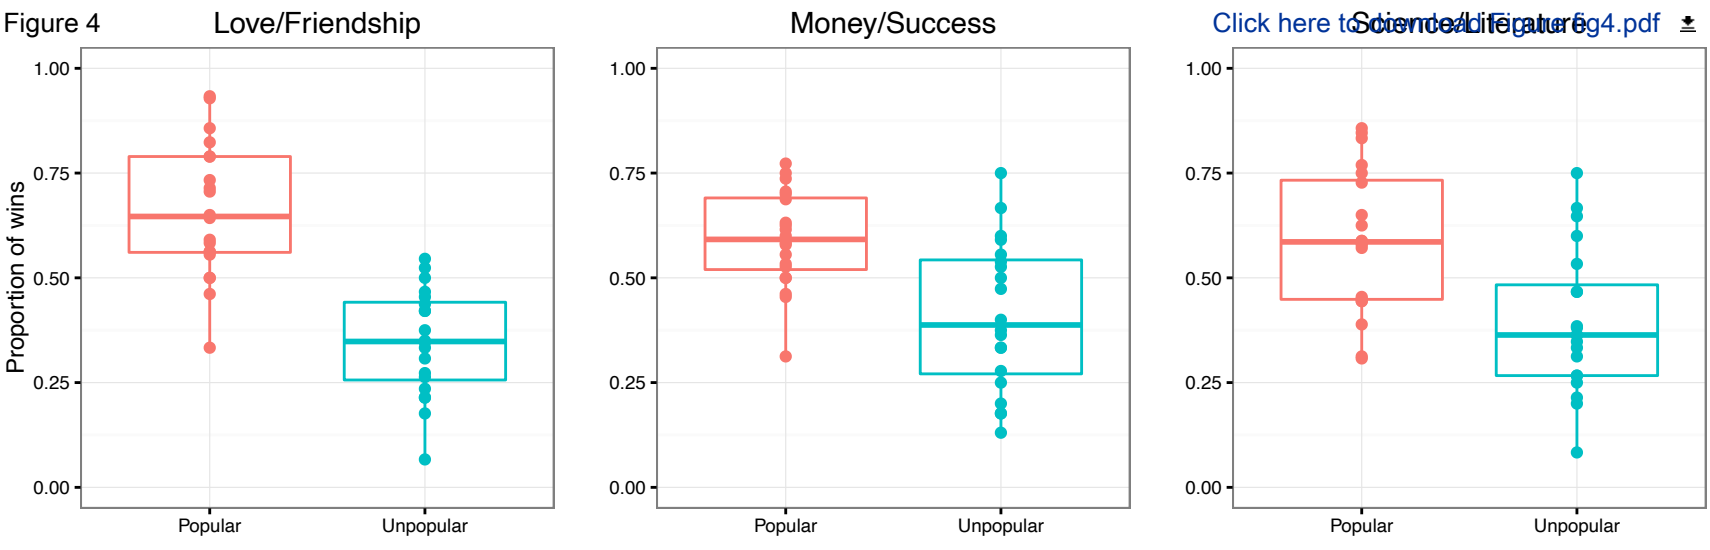

Click here to download Figure fig4.pdf

Figure 5

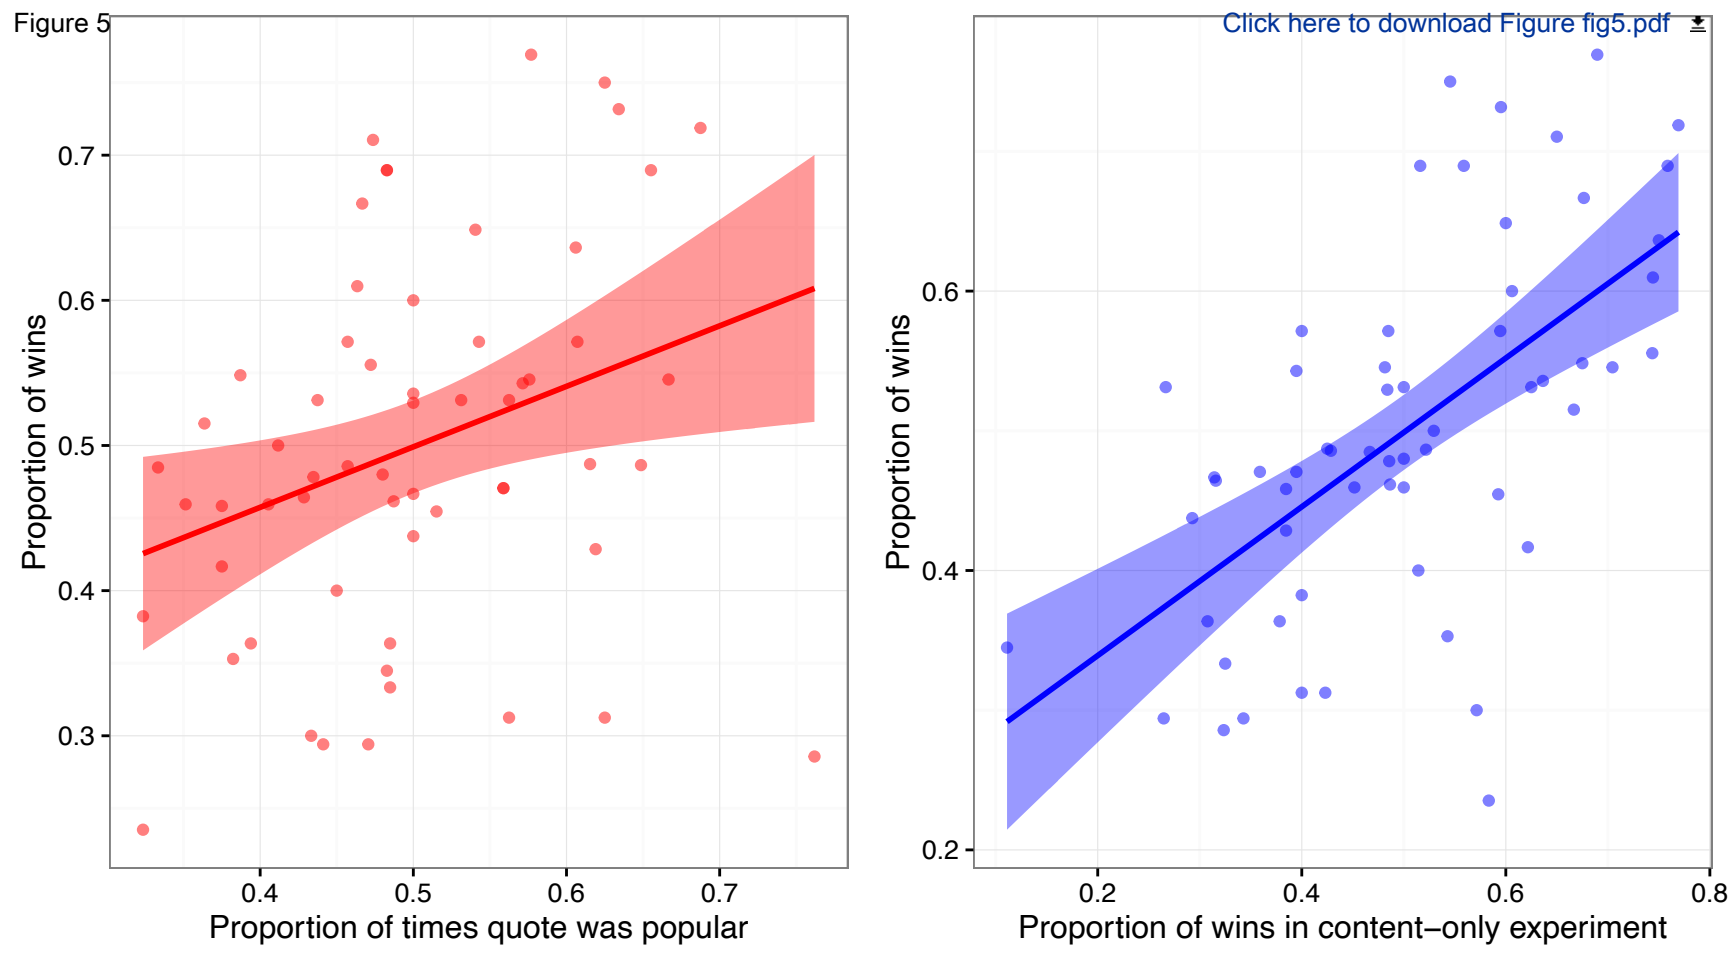
Click here to download Figure fig5.pdf
Proportion of wins
0.7
0.6
0.5
0.4
0.3
0.4
0.5
0.6
0.7
Proportion of times quote was popular
Proportion of wins
0.6
0.4
0.2
0.2
0.4
0.6
0.8
Proportion of wins in content−only experiment

Figure 6

Click here to download Figure fig6.pdf