Bayesian reanalysis of Mesoudi et al. (2015)

Alex Mesoudi

22 May 2016

1 Introduction

This is a reanalysis of the data presented in a previous paper (Mesoudi, Chang, Murray, & Lu, 2015), using methods from Richard McElreath's Statistical Rethinking book (McElreath, 2016) and `rethinking` package (McElreath, 2015). The reanalysis generates virtually identical results to the original analysis, but is hopefully more useful and transparent to other researchers.

In Mesoudi et al. (2015) we found higher rates of social learning in a computer-based artifact-design task amongst participants from mainland China, compared to participants from the UK, from Hong Kong, and Chinese immigrant students studying in the UK. Participants were presented with 29 opportunities to either copy or not copy others during a season of 30 hunts (participants could not copy on the first hunt). Participants could see others' success rates, making social learning payoff-biased and therefore highly beneficial in this challenging task. There were 3 seasons, each with 29 opportunities to copy. Seasons 1 and 2 featured different environments but no within-season environmental change. Season 3 featured a different environment *and* within-season environmental change.

2 Visualisation of data

The original paper presented bar charts of the proportion of copying per season as in Figure 1 below (Figure 2 in Mesoudi et al. (2015)). We can see the higher copying proportion in Chinese Mainland (CM) participants in Seasons 1 and 2, but not in Season 3. In Season 3 the UK, HK and CI participants increased their copying frequency almost up to CM levels.

However, bar charts can obscure variation in data. Ideally visualisations would include representations of the actual data, and counts rather than proportions. Figures 2-4 show some alternatives. These show that there is a broad spread of copying frequencies within each cultural group, but that in Seasons 1 and 2 there are more high-copying CM participants compared to the other three groups.

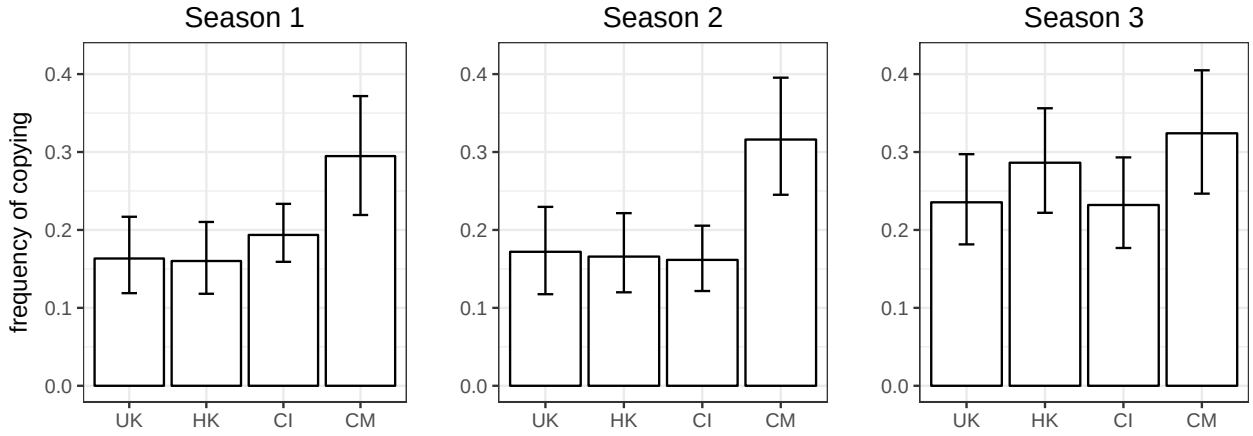

Figure 1: Bar charts as in Mesoudi et al. (2015). UK=British, HK=Hong Kong, CI=Chinese Immigrants, CM=Chinese Mainland.

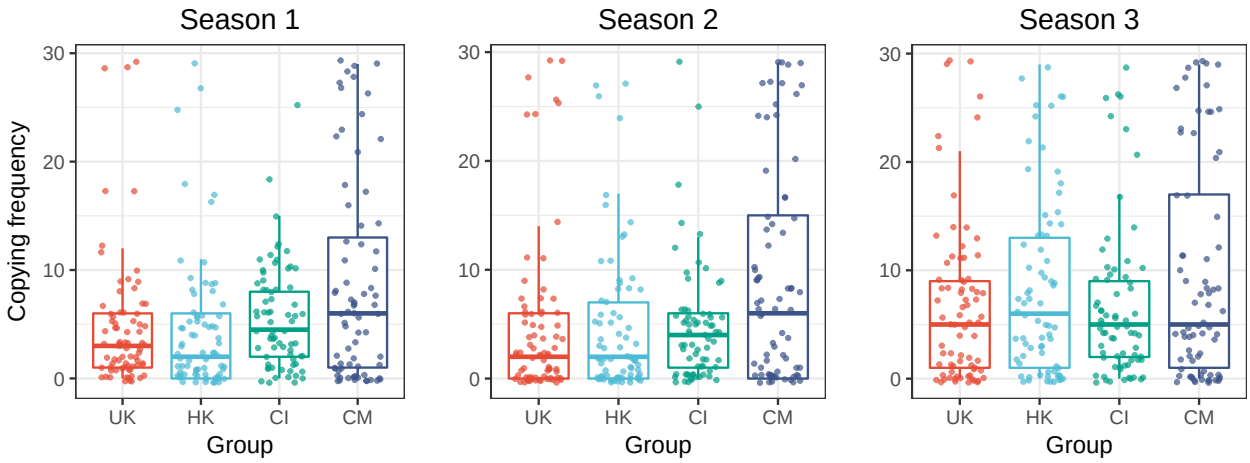

Figure 2: Jitter and box plots

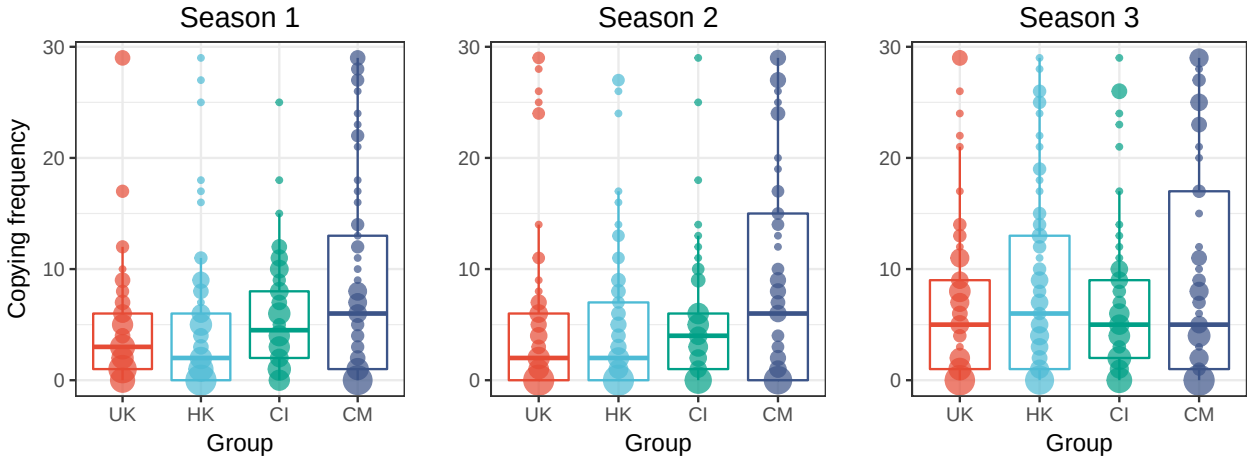

Figure 3: Count and box plots. The size of the circle indicates the number of participants at that value.

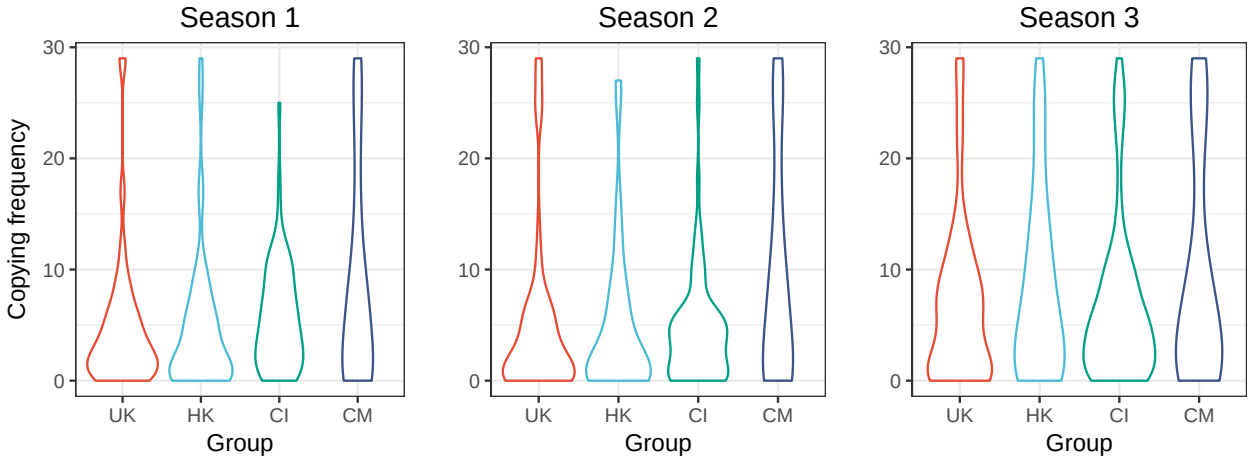

Figure 4: Violin plots, showing narrower bases for CM in Seasons 1 and 2, compared to other groups which are bottom-heavy.

3 Bayesian re-analysis

The original paper used negative binomial regression to compare copying proportions across cultural groups. These regressions are shown in Table 1 of Mesoudi et al. (2015). Here I use the `map` function from the `rethinking` package (McElreath, 2015) to compare cultural groups in an aggregated binomial model. The dependent measure is the number of hunts (out of 29) on which a participant copied another participant. This is performed separately for season 1, 2 and 3, each of which featured 29 opportunities to copy[1]. For each season three models are compared: an intercept-only model, a model with cultural group as a predictor, and a full model with cultural group, age and sex as predictors. UK is the reference group. The code for the full model of Season 1 is:

```
s1.full.model <- map(
  alist(
    season1_copies ~ dbinom( 29 , p ) ,
    logit(p) <- a + bHK*cultureHK + bCI*cultureCI + bCM*cultureCM + bA*age + bF*female,
    a ~ dnorm(0,1) ,
    bHK ~ dnorm(0,1) ,
    bCI ~ dnorm(0,1) ,
    bCM ~ dnorm(0,1) ,
    bA ~ dnorm(0,0.5) ,
    bF ~ dnorm(0,0.5)
  ) ,
  data=d )
```

with Season 2 and 3 identical except for the dependent measures. I also ran these models with MCMC using `map2stan` but the results were virtually identical, so I do not report this here.

3.1 Seasons 1 and 2

For both Seasons 1 and 2 the full model is overwhelmingly best supported:

```
                   WAIC pWAIC dWAIC weight     SE   dSE
s1.full.model    8350.5   5.9   0.0      1 102.19    NA
s1.culture.model 8395.4   4.1  44.9      0 102.06 13.99
s1.null.model    8538.7   1.1 188.2      0 101.24 28.57
```

```
                   WAIC pWAIC dWAIC weight     SE   dSE
s2.full.model    8336.5   6.1   0.0      1 102.04    NA
s2.culture.model 8369.8   4.0  33.3      0 102.08 12.00
s2.null.model    8568.6   1.0 232.1      0 100.99 31.31
```

The full model for Season 1 is shown below, along with exponentiated coefficients to give relative odds (e.g. CM participants are 2.24 times more likely to copy compared to UK participants). HK participants are virtually identical in copying frequency to UK participants. CI are slightly higher, and CM are much higher. Women copy more than men, and age has little effect.

[1] I do not run a single regression with season as a within-participant factor because, as noted above, Season 3 introduces within-season environmental change and so is not comparable to the others. Seasons 1 and 2 could be combined, but the gain in having one fewer model does not seem to me to outweigh the cost of having coefficients that are harder to interpret.

```
     Mean StdDev  5.5% 94.5%
a   -0.89   0.29 -1.36 -0.42
bHK -0.01   0.08 -0.14  0.12
bCI  0.16   0.08  0.03  0.29
bCM  0.80   0.08  0.68  0.93
bA  -0.05   0.01 -0.07 -0.02
bF   0.32   0.06  0.23  0.41
```

```
        a       bHK       bCI       bCM        bA        bF
0.4100784 0.9891871 1.1722992 2.2358069 0.9558442 1.3749853
```

The full model for Season 2 is shown below and is almost the same as for Season 1. HK participants are again virtually identical to UK participants, CI slightly lower, and CM much higher. Women again copy more than men, and age again has little effect.

```
     Mean StdDev  5.5% 94.5%
a   -1.37   0.29 -1.83 -0.91
bHK -0.03   0.08 -0.16  0.10
bCI -0.13   0.08 -0.26  0.00
bCM  0.82   0.07  0.70  0.94
bA  -0.02   0.01 -0.04  0.00
bF   0.32   0.06  0.23  0.41
```

```
        a       bHK       bCI       bCM        bA        bF
0.2533312 0.9687078 0.8797198 2.2641814 0.9819492 1.3771342
```

3.2 Season 3

For Season 3, the full model and culture model both received support but the full model slightly more:

```
                   WAIC pWAIC dWAIC weight    SE   dSE
s3.full.model    9814.4   6.0   0.0   0.63 82.28    NA
s3.culture.model 9815.5   4.1   1.1   0.37 82.30  4.77
s3.null.model    9871.6   1.1  57.2   0.00 81.41 16.40
```

The full model is shown below. CI participants are virtually identical to UK participants, while HK and CM are higher. Both age and sex have very weak effects.

```
     Mean StdDev  5.5% 94.5%
a   -0.62   0.26 -1.03 -0.21
bHK  0.26   0.07  0.15  0.37
bCI -0.02   0.07 -0.14  0.09
bCM  0.46   0.07  0.35  0.57
bA  -0.03   0.01 -0.05 -0.01
bF   0.02   0.05 -0.06  0.10
```

```
        a       bHK       bCI       bCM        bA        bF
0.5389395 1.2955908 0.9778064 1.5860515 0.9725006 1.0202826
```

4 Summary

The coefficients and confidence intervals shown here are very similar to the original regression model results shown in Mesoudi et al. (2015), but hopefully more straightforward and understandable. I will soon incorporate task performance into the above, to complete the reanalysis of the original paper.

The RMarkdown file containing code for running these reanalyses and producing this document is available at:

https://alex-mesoudi.squarespace.com/s/procb_reanalysis_22may16.Rmd

References

McElreath, R. (2015). Rethinking: Statistical Rethinking Book Package.

McElreath, R. (2016). *Statistical rethinking: A Bayesian course with examples in r and stan*. CRC Press.

Mesoudi, A., Chang, L., Murray, K., & Lu, H. (2015). Higher frequency of social learning in China than in the West shows cultural variation in the dynamics of cultural evolution. *Proceedings of the Royal Society B*, *282*, 20142209.