SUPPLEMENTARY MATERIAL

For

MORAL THIN-SLICING:

FORMING MORAL IMPRESSIONS FROM A BRIEF GLANCE

Julian De Freitas[1*†] & Alon Hafri[2*]

1 – Assistant Professor, Harvard Business School

2 – Assistant Professor, Departments of Linguistics and Cognitive Science, University of *Delaware*

* These authors contributed equally to this work.

† Corresponding author: jdefreitas@hbs.edu

Contents:

Pre-Study: Survey on Frequency of Witnessing Events in Different Contexts

Our intuitions were that, on the one hand, it is uncommon these days to witness moral transgressions firsthand — at least ones involving physical harm — relative to past eras; while on the other hand, video and social media have made moral transgressions more accessible visually than ever before. We sought to empirically check these intuitions. To do so, we ran an online study in which we asked participants to report the last time they visually witnessed our different social interaction categories (*hugging, tickling, slapping,* etc.) firsthand or online.

Method

We recruited 30 participants from Prolific. Participants were told to report when they last witnessed different types of events in two different contexts: "in-person" (first-hand) and "online or in the news" (on the internet, on an app, or on TV). On each trial, participants viewed a description of one of the 27 social interaction categories from our main studies, in the following form: "One person Xing another" (where "X" was filled in with the category label, e.g., "slap"). Participants were tasked with choosing from among the following options (with ordinal response codings for analyses listed in parentheses): "never" (0), "more than a year ago" (1), "within the last year" (2), "within the last month" (3), "within the last week" (4), and "today" (5). Participants responded separately about in-person and online contexts on the same trial. We grouped results separately into low- and high-harm social interaction categories according to the categorizations from Study S1 (below).

An additional attention-check trial was included, in which the phrase presented was *This is an attention-check trial; Select the "Attention Check" response options below.* For this trial, participants were told to select the option "This is an attention check trial." One participant was excluded for selecting the attention-check option on non-attention-check trials.

Results

Results are shown in Figure S1. Participants reported that low-harm social interactions were last visually experienced on average close to within the last month, both in firsthand (M = 2.82, *Median* = 3.0, SD = 1.46) and online (M = 2.81, *Median* = 3.0, SD = 1.31) contexts. A paired t-test with Context as a factor showed no significant difference between these conditions, $t(28) = 0.09$, $p = .926$, $d = 0.02$.

By contrast, for high-harm social interactions, there was a marked difference by context: participants reported that high-harm social interactions were last visually experienced on average close to more than a year ago for firsthand contexts (M = 0.92, *Median* = 1.0, SD = 1.06) but within the last year for online contexts (M = 2.40, *Median* = 2.0, SD = 1.18). A paired t-test with Context as a factor showed a significant difference between these two conditions, $t(28) = 13.41$, $p < .001$, $d = 2.49$. Furthermore, a repeated measures ANOVA with Context and Harm Level as factors showed a significant interaction, $F(1,28) = 174$, $p < .001$, $\eta_p^2 = 0.86$.

These results confirmed our intuitions that visually witnessing such social interactions in-person is relatively rare, but in an online context, it is relatively common.

RESULTS FROM PRE-STUDY SURVEY

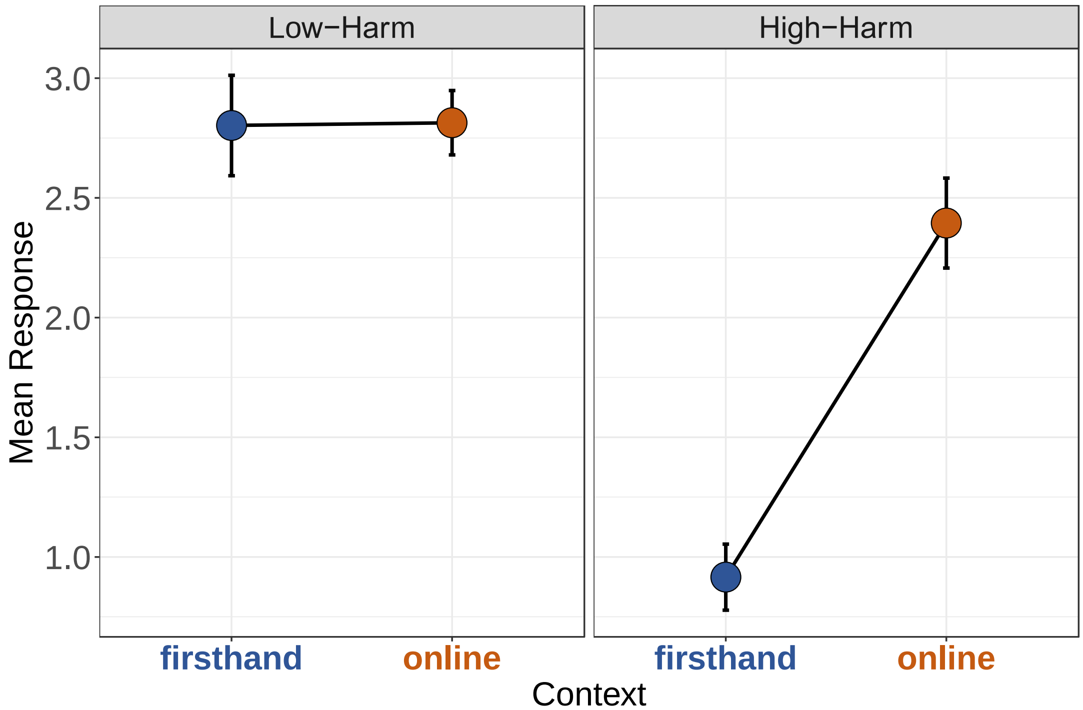

Figure S1. Mean responses for when participants last visually witnessed low-harm social interactions (left panel) and high-harm interactions (right panel), in firsthand and online contexts. Responses were coded ordinally, from 0 (“never”) to 5 (“today”). Points are means across participants; error bars are 95% confidence intervals, correcting the variance for repeated measures across participants.

Study S1: Controlled Image Norming

Since we ultimately sought to determine how consistently participants extract color, role, harm, and moral wrongness from brief image presentations (i.e., whether they do so in ways similar to how they do so under no viewing-time constraints), we needed to include only images in the main studies that observers ordinarily categorized similarly when given unlimited viewing time. To this end, we first ran an *unspeeded* norming study in which a subsequent trial was only presented once the observer finished responding to the current trial.

Method

We recruited 169 participants from Amazon's Mechanical Turk and excluded 14 using the same exclusion criteria specified in Study 1 (apart from constraints on web browser display time). This left 155 participants for analysis (97 identifying as male, 58 as female; mean age 35.7, *sd* 10.7, range 19–73). Stimuli were the same as used in Study 1 of the main manuscript, i.e., 108 images of identical-twin actors engaged in various social interaction categories varying in the degree of harm, with identity and spatial location of the Agent and Patient fully counterbalanced across categories. The design was identical to Study 1 in the main manuscript, except that here, the image and probe appeared at the same time, and both remained on screen until response. The sample sizes for each condition after exclusions were the following: Color (n = 38), Role (n = 24), Harm (n = 30), and Moral Wrongness (n = 63).

Results

We excluded any social interaction category with one of the agent exemplars (red-Agent or blue-Agent) more than 2.5 SDs below each task's mean agreement, collapsing over spatial locations of the agent (left or right). By "agreement", we mean the extent to which the majority of participants answered the same way (e.g., that the action was harmful). These thresholds were 89.6% for color, 63.9% for role, 68.7% for harm, and 68.9% for moral wrongness. Three social interaction categories — *scaring*, *listening to*, and *pulling* — had low average agreement on at least one of the four tasks, so were excluded (*scaring* blue-Agent: 89.5% color agreement; *listening to* blue-Agent: 54.2% role agreement; *listening to* red-Agent: 39.6% role agreement; *pulling* blue-Agent: 59.3% harm agreement, 68.3% moral wrongness agreement; *pulling* red-Agent: 66.7% harm agreement, 66.7% moral wrongness agreement). The remaining social interaction categories had high average agreement rates across all tasks (Color: 96%; Role: 89%; Harm: 92%; Moral Wrongness (agent only): 90%), so were all used in subsequent studies.

Note that in our primary analyses for studies reported in the main manuscript, we only report moral wrongness values for the agent rather than patient because in harmful social interactions the agent is the one who is considered most morally wrong (Gray & Wegner, 2009). In line with this assumption, patients across all social interaction categories were rated as lower in moral wrongness than agents (agents: M=0.39, SD=0.40, range 0.02–0.95; patients: M=0.12, SD=0.11, range 0.02–0.37; agent-patient difference: M=0.26, SD=0.29, range -0.01–0.71). This was also the case for both unharmful interactions (agents: M=0.09, SD=0.05, range 0.02–0.18; patients: M=0.04, SD=0.02, range 0.02–0.10; agent-patient difference: M=0.04, SD=0.04, range -0.01–0.12) and harmful ones (agents: M=0.88, SD=0.06, range 0.77–0.95; patients: M=0.25, SD=0.07, range 0.19–0.37; agent-patient difference: M=0.63, SD=0.06, range 0.55–0.71).

Discussion

This norming study curated a set of social interaction images for which participants provided similar judgments of color, role, harm, and moral wrongness, given no time constraints. In the subsequent studies, we investigate whether participants can make judgments consistent with these unspeeded responses under speeded presentation.

Table S1

RESULTS OF TWO-SAMPLE BAYES FACTOR *T*-TESTS IN STUDY 2

Dur. (ms)	**Color**	**Role**	**Harm**	**Moral Wrongness**
17	n=26, BF=0.01, $t(37.58)=3.55^{*}$	n=19, **BF=1.32,** $t(40.72)=1.59^{ns}$	n=27, BF=0, $t(53.33)=8.65^{*}$	n=38, BF=0, $t(96.57)=11.97^{*}$
33	n=27, BF=0.03, $t(50.61)=3.37^{*}$	n=26, BF=0.5, $t(36.66)=2.17^{ns}$	n=28, BF=0, $t(53.85)=6.58^{*}$	n=44, BF=0, $t(102.44)=9.43^{*}$
50	n=30, BF=0.04, $t(56.86)=3.31^{*}$	n=23, BF=0.87, $t(44.9)=1.86^{ns}$	n=34, BF=0, $t(53.8)=6.43^{*}$	n=55, BF=0, $t(111.03)=9.43^{*}$
67	n=27, BF=0.14, $t(39.01)=2.64^{ns}$	n=28, **BF=2.1,** $t(44.5)=1.13^{ns}$	n=30, BF=0, $t(56.54)=4.5^{*}$	n=51, BF=0, $t(110.37)=9.02^{*}$
83	n=25, BF=0.01, $t(39.64)=3.6^{*}$	n=29, **BF=2.07,** $t(44.78)=1.14^{ns}$	n=28, BF=0, $t(55.32)=5.95^{*}$	n=51, BF=0, $t(106.23)=8.85^{*}$
100	n=27, BF=0.94, $t(46.33)=1.79^{ns}$	n=34, BF=0.84, $t(41.02)=1.82^{ns}$	n=36, BF=0, $t(58.34)=5.38^{*}$	n=55, BF=0, $t(104.69)=9.28^{*}$
133	n=31, BF=0.24, $t(65.17)=2.61^{ns}$	n=30, BF=0.55, $t(50.33)=2.18^{ns}$	n=29, BF=0, $t(55.64)=4.32^{*}$	n=57, BF=0, $t(111.39)=8.64^{*}$
150	n=30, BF=0.55, $t(65.6)=2.21^{ns}$	n=25, **BF=2.84,** $t(45.69)=0.71^{ns}$	n=28, BF=0, $t(55.98)=4.97^{*}$	n=50, BF=0, $t(108.78)=8.35^{*}$
167	n=31, BF=0.02, $t(57.0)=3.55^{*}$	n=20, **BF=2.34,** $t(40.76)=0.95^{ns}$	n=31, BF=0, $t(56.8)=5.64^{*}$	n=55, BF=0, $t(108.72)=7.74^{*}$
200	n=32, BF=0.92, $t(65.51)=1.87^{ns}$	n=30, **BF=2.99,** $t(45.04)=0.67^{ns}$	n=35, BF=0, $t(58.52)=4.26^{*}$	n=55, BF=0, $t(115.99)=6.12^{*}$
250	n=31, BF=0.45, $t(65.94)=2.3^{ns}$	n=32, **BF=3.62,** $t(50.19)=0.2^{ns}$	n=31, BF=0.08, $t(58.93)=3.12^{*}$	n=42, BF=0, $t(99.15)=5.26^{*}$
500	n=33, BF=0.89, $t(65.61)=1.88^{ns}$	n=31, **BF=3.37,** $t(47.99)=-0.43^{ns}$	n=25, **BF=3.21,** $t(51.76)=0.57^{ns}$	n=32, **BF=1.17,** $t(80.3)=1.93^{ns}$
750	n=25, BF=0.63, $t(48.76)=2.06^{ns}$	n=34, **BF=2.21,** $t(38.89)=-1.05^{ns}$	n=25, **BF=2.03,** $t(51.52)=-1.2^{ns}$	n=27, **BF=3.56,** $t(54.15)=0.64^{ns}$
1000	n=32, **BF=3.01,** $t(67.81)=0.84^{ns}$	n=31, **BF=3.44,** $t(34.31)=-0.34^{ns}$	n=28, **BF=1.51,** $t(53)=-1.5^{ns}$	n=30, **BF=4.33,** $t(73.48)=0^{ns}$
1500	n=37, **BF=3.59,** $t(72.64)=0.6^{ns}$	n=32, **BF=3.67,** $t(47.13)=-0.01^{ns}$	n=25, BF=0.99, $t(51.5)=-1.84^{ns}$	n=41, **BF=2.01,** $t(96.66)=-1.58^{ns}$

Bayes factor *t*-tests compared each speeded judgment condition to the unspeeded condition of Study 2: values above 1 indicate evidence in favor of the null (no difference) over the alternative (BF_{01}, shortened to BF here). Bolded cells indicate more evidence for the null than the alternative. Results of standard two-sample *t*-tests are also given for illustrative purposes, although this was not the primary measure of interest for this study. * $p < .05$, *ns:* $p > .05$. Significance of p values is based on correction for the fifteen comparisons conducted within each task, using the Bonferroni-Holm method. Dur. = Duration.

Response Biases Across Tasks (Studies 3 and 4)

As in Study 1, we wanted to explore whether general tendencies to respond “yes” or “no” would vary by task in ways that could reveal how strategies differ in this new stimulus set. To explore this question, we calculated criterion or *c* for each participant (current study only) and then tested significance of *c* values across participants relative to zero (chance), separately for each task. We did so separately for Studies 3 and 4, with results appearing below.

Study 3 Response Bias Analysis

Similarly to Study 1, participants in the Role task did not exhibit any significant response biases, as evinced by criterion values that were not significantly different from zero (M = –0.07, 95% CI [–0.18, 0.04], $t(26) = 1.28$, $p = .213$, $d = 0.07$); and participants in the Harm and Moral Wrongness tasks again showed a significant bias to respond “yes”, as evinced by significant negative *c* values (Harm: M = –0.43, 95% CI [–0.60, –0.26], $t(28) = 5.14$, $p < .001$, $d = 0.95$; Moral Wrongness: M = –0.23, 95% CI [–0.44, –0.01], $t(33) = 2.14$, $p = .040$, $d = 0.37$). Additionally, participants in the Color task showed a bias to respond “no”, as evinced by significant positive *c* values ($M = 0.10$, 95% CI [0.02, 0.18], $t(28) = 2.61$, $p = .014$, $d = 0.49$).

Study 4 Response Bias Analysis

Participants in the Color task were conservative in their responding, with a significant bias to respond “no” on the whole ($M = 0.85$, 95% CI [0.46, 1.24], $t(20) = 4.57$, $p < .001$, d = 1.00). However, the other three tasks did not exhibit any significant response biases, as evinced by criterion values that were not significantly different from zero (Role: $M = 0.21$, 95% CI [–0.04, 0.47], $t(22) = 1.74$, $p = .096$, $d = 0.36$; Harm: $M = 0.02$, 95% CI [–0.40, 0.44], $t(24) = 0.12$, *p* = .904, *d* = 0.02; Moral Wrongness: *M* = –0.08, 95% CI [–0.37, 0.21], *t*(44) = 0.59, *p* = .561, *d* = 0.09).

The lack of response biases for harm and moral wrongness in the current study is especially notable in contrast to the significant “yes” biases observed in Studies 1 and 3. Although we should interpret the differences in response biases across experiments with caution, we speculate that participants in the current study were aware of the difficulty of extracting fine-grained features useful for categorizing harm level in the darkened images; and that participants in Studies 1 and 3, by contrast, may have (falsely) believed they were extracting information indicative of harm.

Mixed-Effects Modeling Analyses for Studies 1, 3, and 4

Throughout the main manuscript, we assessed the agreement between speeded and unspeeded responses by computing the reliability of *d'* across participants. However, we find comparable results with analyses of individual trial-level 'match' responses (hits and correct rejections) using mixed-effects logistic regression models. While these models do not offer bias-free performance measures, they do offer complementary advantages to our signal-detection-theory approach: they enable generalization of statistical inferences simultaneously across participants and items (social interaction categories) by accounting for both participant- and item-level variability (Baayen, Davidson, & Bates, 2008; Barr, Levy, Scheepers, & Tily, 2013), and they also deal well with missing trials and unbalanced data.

We report results of these analyses below for Studies 1, 3, and 4. For Study 1, we report separate analyses for each task (i.e., Color, Role, Harm, and Moral Wrongness). For Study 3, we report comparisons with Study 1, separately for each task. For Study 4, we report comparisons both with Study 1 and with Study 3, separately for each task.

Analyses

The dependent variable in each study was whether a response was a 'match' or not. (Matches encompassed both hits and correct rejections, according to the binary categorizations from the unspeeded norming results in Study S1.) In all three studies, the independent variable Present (`present`) indicated whether the response (according to unspeeded norming) should be "yes" (`present == 1`, for which a 'match' response was a hit) or "no" (`present == 0`, for which a 'match' response was a correct rejection). In Studies 3 and 4, the variable of interest was Experiment Name (`expName`) — whether there was a difference between the data from the two studies being compared. In Studies 3 and 4, the variable Harm Level (harmLevel) indicated whether the social interaction category was high-harm (harmLevel == 1) or low-harm (harmLevel == 1); this variable was also tested for significance. All of the independent variables above were binary and were sum-coded as −1.0 and 1.0.

In all cases, we tested for significance of variables by using likelihood-ratio tests on the χ^2 values from nested model comparisons with the same random-effects structure. We started with the maximal random-effects structure: correlated random intercepts and random slopes by participant (part_id) and by item (social interaction category, i.e., action). If models did not converge, we simplified the random effects structure by first using uncorrelated intercepts and slopes, and we followed that by dropping random intercepts and slopes until convergence, starting with those that accounted for the least variance.

Analyses were conducted in R (version 4.0.2) with the *lme4* package (version 1.1-34). Models are specified below in terms of R formula syntax.

Results for Study 1

Note: A significant intercept under Fixed effects indicates that ‘match’ responses were significantly above chance (50%) for that task.

Color:

```
Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
 Family: binomial  ( logit )
Formula: match ~ 1 + (1 + present | part_id)
   Data: subset(tmpdata, speedCond == "speeded" & judgmentType == "color")

     AIC      BIC   logLik deviance df.resid
  1105.5   1128.3   -548.8   1097.5     2173

Scaled residuals:
    Min      1Q  Median      3Q     Max
```

```
-6.5701  0.1243  0.2183  0.2944  3.4694

Random effects:
 Groups   Name        Variance Std.Dev. Corr 
 part_id  (Intercept) 2.6328   1.6226        
          present1    0.2548   0.5048   -0.38
Number of obs: 2177, groups:  part_id, 23

Fixed effects:
            Estimate Std. Error z value Pr(>|z|)    
(Intercept)   2.9204     0.3492   8.364   <2e-16 ***
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
```

Role:

```
Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
 Family: binomial  ( logit )
Formula: match ~ 1 + (1 + present | part_id) + (1 | action)
   Data: subset(tmpdata, speedCond == "speeded" & judgmentType == "role")

     AIC      BIC   logLik deviance df.resid 
  1641.0   1668.7   -815.5   1631.0     1894 

Scaled residuals: 
    Min      1Q  Median      3Q     Max 
-5.8069  0.2196  0.3331  0.4689  1.1971 

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 action   (Intercept) 0.1013   0.3183       
 part_id  (Intercept) 0.7638   0.8739       
          present1    0.1016   0.3188   0.53
Number of obs: 1899, groups:  action, 24; part_id, 20

Fixed effects:
            Estimate Std. Error z value Pr(>|z|)    
(Intercept)   1.8518     0.2061   8.986   <2e-16 ***
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
```

Harm:

```
Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
 Family: binomial  ( logit )
Formula: match ~ (1 + present | part_id) + (1 | action) + present
   Data: subset(tmpdata, speedCond == "speeded" & judgmentType == "harm")
```

```
   AIC      BIC   logLik deviance df.resid 
  1922.3   1955.8   -955.1   1910.3     1970 

Scaled residuals: 
    Min      1Q  Median      3Q     Max 
-6.5708 -0.5154  0.3247  0.4972  2.7262 

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 action   (Intercept) 0.8194   0.9052       
 part_id  (Intercept) 0.4943   0.7031       
          present1    0.4491   0.6701   0.34
Number of obs: 1976, groups:  action, 24; part_id, 21

Fixed effects:
            Estimate Std. Error z value Pr(>|z|)    
(Intercept)   1.4314     0.2562   5.588  2.3e-08 ***
present1     -0.5718     0.2517  -2.272   0.0231 *  
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Correlation of Fixed Effects:
         (Intr)
present1 -0.067
```

Moral Wrongness:

```
Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
 Family: binomial  ( logit )
Formula: match ~ 1 + (1 + present | part_id) + (1 | action)
   Data: subset(tmpdata, speedCond == "speeded" & judgmentType == "moral.Agent")

   AIC      BIC   logLik deviance df.resid 
  2029.1   2056.9  -1009.5   2019.1     1927 

Scaled residuals: 
    Min      1Q  Median      3Q     Max 
-5.3803 -0.6272  0.3340  0.5645  3.3921 

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 part_id  (Intercept) 0.3638   0.6032       
          present1    0.9222   0.9603   0.62
 action   (Intercept) 0.7881   0.8877       
Number of obs: 1932, groups:  part_id, 41; action, 24

Fixed effects:
            Estimate Std. Error z value Pr(>|z|)    
(Intercept)   1.1715     0.2217   5.283 1.27e-07 ***
```

```
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
```

Results for Comparisons of Study 3 and Study 1

Color:

```
Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
 Family: binomial  ( logit )
Formula: match ~ (1 + present | part_id) + (1 + present | action) + present
   Data: subset(tmpdata, speedCond == "speeded" & judgmentType == "color")

     AIC      BIC   logLik deviance df.resid
  7312.4   7372.8  -3648.2   7296.4    14126

Scaled residuals:
    Min      1Q  Median      3Q     Max
-9.4386  0.1290  0.2298  0.3271  3.2009

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 part_id  (Intercept) 1.63944  1.2804
          present1    0.29886  0.5467   0.09
 action   (Intercept) 0.15022  0.3876
          present1    0.05883  0.2426   0.22
Number of obs: 14134, groups:  part_id, 52; action, 22

Fixed effects:
            Estimate Std. Error z value Pr(>|z|)
(Intercept)   2.9002     0.2033  14.269   <2e-16 ***
present1      0.2649     0.1052   2.519   0.0118 *
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Correlation of Fixed Effects:
         (Intr)
present1 0.129
```

Role:

```
Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
 Family: binomial  ( logit )
Formula: match ~ (1 + present + harmLevel | part_id) + (1 + present + expName |      action) + harmLevel +
expName
   Data: subset(tmpdata, speedCond == "speeded" & judgmentType == "role")
```

```
    AIC      BIC   logLik deviance df.resid 
 11521.0  11632.9  -5745.5  11491.0    12853 

Scaled residuals: 
    Min      1Q  Median      3Q     Max 
-5.7527  0.1430  0.3274  0.4996  3.9332 

Random effects:
 Groups   Name        Variance Std.Dev. Corr       
 part_id  (Intercept) 0.88282  0.9396              
          present1    0.23360  0.4833    0.11      
          harmLevel1  0.16151  0.4019   -0.05  0.00
 action   (Intercept) 0.18487  0.4300              
          present1    0.15941  0.3993   -0.70      
          expName1    0.02781  0.1668   -0.35  0.32
Number of obs: 12868, groups:  part_id, 47; action, 22

Fixed effects:
            Estimate Std. Error z value Pr(>|z|)    
(Intercept)   1.7524     0.1685  10.401   <2e-16 ***
harmLevel1   -0.2336     0.1022  -2.285   0.0223 *  
expName1      0.2914     0.1486   1.961   0.0499 *  
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Correlation of Fixed Effects:
           (Intr) hrmLv1
harmLevel1 -0.084       
expName1    0.103 -0.022
```

Harm:

First, model comparisons revealing no effect of Experiment Name:

```
> anova(lm.0, lm.1, lm.2, lm.3)
Data: subset(tmpdata, speedCond == "speeded" & judgmentType == "harm")
Models:
lm.0: match ~ 1 + (1 + present | part_id) + (1 + expName | action)
lm.1: match ~ (1 + present | part_id) + (1 + expName | action) + present
lm.2: match ~ (1 + present | part_id) + (1 + expName | action) + present + expName
lm.3: match ~ (1 + present | part_id) + (1 + expName | action) + present + expName + present:expName
     npar   AIC   BIC  logLik deviance  Chisq Df Pr(>Chisq)   
lm.0    7 12969 13022 -6477.6    12955                        
lm.1    8 12963 13023 -6473.3    12947 8.5287  1   0.003496 **
lm.2    9 12964 13032 -6472.9    12946 0.7277  1   0.393635   
lm.3   10 12966 13041 -6472.8    12946 0.3547  1   0.551458   
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
```

```
> anova(lm.0, lm.1, lm.3)
Data: subset(tmpdata, speedCond == "speeded" & judgmentType == "harm")
Models:
lm.0: match ~ 1 + (1 + present | part_id) + (1 + expName | action)
lm.1: match ~ (1 + present | part_id) + (1 + expName | action) + present
lm.3: match ~ (1 + present | part_id) + (1 + expName | action) + present + expName + present:expName
     npar   AIC   BIC  logLik deviance  Chisq Df Pr(>Chisq)
lm.0    7 12969 13022 -6477.6    12955
lm.1    8 12963 13023 -6473.3    12947 8.5287  1   0.003496 **
lm.3   10 12966 13041 -6472.8    12946 1.0824  2   0.582050
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Now, the best-fitting model:

```
Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
 Family: binomial  ( logit )
Formula: match ~ (1 + present | part_id) + (1 + expName | action) + present
   Data: subset(tmpdata, speedCond == "speeded" & judgmentType == "harm")

     AIC      BIC   logLik deviance df.resid
 12962.6  13022.8  -6473.3  12946.6    13732

Scaled residuals:
    Min      1Q  Median      3Q     Max
-9.4553 -0.5507  0.2952  0.5378  4.0835

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 part_id  (Intercept) 0.4532   0.6732
          present1    0.7601   0.8718   -0.01
 action   (Intercept) 0.7509   0.8665
          expName1    0.0788   0.2807   -0.07
Number of obs: 13740, groups:  part_id, 50; action, 22

Fixed effects:
            Estimate Std. Error z value Pr(>|z|)
(Intercept)   1.3718     0.2143   6.401 1.54e-10 ***
present1     -0.7046     0.2278  -3.093  0.00198 **
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Correlation of Fixed Effects:
         (Intr)
present1 -0.143
```

Moral Wrongness:

First, model comparisons revealing a marginal effect of Experiment Name:

```
> anova(lm.0, lm.1, lm.2, lm.3)
Data: subset(tmpdata, speedCond == "speeded" & judgmentType == "moral.Agent")
Models:
lm.0: match ~ 1 + (1 + present | part_id) + (1 + expName | action)
lm.1: match ~ (1 + present | part_id) + (1 + expName | action) + present
lm.2: match ~ (1 + present | part_id) + (1 + expName | action) + present + expName
lm.3: match ~ (1 + present | part_id) + (1 + expName | action) + present + expName + present:expName
     npar    AIC    BIC  logLik deviance  Chisq Df Pr(>Chisq)
lm.0    7 9655.4 9705.5 -4820.7   9641.4
lm.1    8 9653.9 9711.2 -4819.0   9637.9 3.4463  1    0.06339 .
lm.2    9 9653.1 9717.5 -4817.5   9635.1 2.8423  1    0.09181 .
lm.3   10 9655.0 9726.6 -4817.5   9635.0 0.0694  1    0.79226
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
```

Now, the marginally better-fitting model with the variable Experiment Name included:

```
Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
 Family: binomial  ( logit )
Formula: match ~ (1 + present | part_id) + (1 + expName | action) + present +      expName
   Data: subset(tmpdata, speedCond == "speeded" & judgmentType == "moral.Agent")

     AIC      BIC   logLik deviance df.resid
  9653.1   9717.5  -4817.5   9635.1     9459

Scaled residuals:
    Min      1Q  Median      3Q     Max
-7.5964 -0.6279  0.3031  0.5837  3.7377

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 part_id  (Intercept) 0.31309  0.5595
          present1    1.22958  1.1089   0.39
 action   (Intercept) 0.61720  0.7856
          expName1    0.07422  0.2724   -0.03
Number of obs: 9468, groups:  part_id, 75; action, 22

Fixed effects:
            Estimate Std. Error z value Pr(>|z|)
(Intercept)  1.04314    0.18509   5.636 1.74e-08 ***
present1    -0.41033    0.21558  -1.903   0.0570 .
expName1     0.15630    0.09167   1.705   0.0882 .
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Correlation of Fixed Effects:
         (Intr) prsnt1
present1   -0.057
expName1   -0.022 -0.007
```

Results for Comparisons of Study 4 and Study 1

Color:

```
Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
 Family: binomial  ( logit )
Formula: match ~ (1 + present + harmLevel | part_id) + present + expName +      present:expName
   Data: subset(tmpdata.1.4, speedCond == "speeded" & judgmentType ==      "color")

     AIC      BIC   logLik deviance df.resid
  7017.8   7089.6  -3498.9   6997.8     9716

Scaled residuals:
    Min      1Q  Median      3Q     Max
-5.5167 -0.0589  0.1527  0.3863  3.7467

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 part_id  (Intercept) 1.61825  1.2721
          present1    3.55973  1.8867   -0.14
          harmLevel1  0.08436  0.2904   -0.22 -0.04
Number of obs: 9726, groups:  part_id, 44

Fixed effects:
                  Estimate Std. Error z value Pr(>|z|)
(Intercept)         1.6954     0.2052   8.263  < 2e-16 ***
present1            1.1648     0.2986   3.901 9.59e-05 ***
expName1            1.2695     0.2467   5.147 2.65e-07 ***
present1:expName1  -1.0052     0.3611  -2.784  0.00537 **
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Correlation of Fixed Effects:
            (Intr) prsnt1 expNm1
present1    -0.109
expName1     0.041 -0.028
prsnt1:xpN1  0.024  0.055 -0.222
```

Role:

```
Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
 Family: binomial  ( logit )
Formula: match ~ (1 + present + harmLevel | part_id) + (1 + present + expName |      action) + expName
   Data: subset(tmpdata.1.4, speedCond == "speeded" & judgmentType ==      "role")

     AIC      BIC   logLik deviance df.resid
  8495.2   8596.1  -4233.6   8467.2     9936

Scaled residuals:
    Min      1Q  Median      3Q     Max
-6.2410  0.0571  0.2947  0.4714  4.2916

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 part_id  (Intercept) 1.01229  1.0061
          present1    1.07032  1.0346    0.05
          harmLevel1  0.11519  0.3394   -0.01 -0.19
 action   (Intercept) 0.17145  0.4141
          present1    0.27701  0.5263   -0.24
          expName1    0.02687  0.1639   -0.36 -0.06
Number of obs: 9950, groups:  part_id, 43; action, 19

Fixed effects:
            Estimate Std. Error z value Pr(>|z|)
(Intercept)   1.7417     0.1878   9.276   <2e-16 ***
expName1      0.3425     0.1707   2.007   0.0447 *
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Correlation of Fixed Effects:
         (Intr)
expName1 0.004
```

Harm:

```
Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
 Family: binomial  ( logit )
Formula: match ~ (1 + present | part_id) + (1 + present + expName | action) +      expName
   Data: subset(tmpdata.1.4, speedCond == "speeded" & judgmentType ==      "harm")

     AIC      BIC   logLik deviance df.resid
  9038.2   9118.1  -4508.1   9016.2    10493

Scaled residuals:
    Min      1Q  Median      3Q     Max
-8.2336 -0.4281  0.2173  0.4937  3.6798

Random effects:
 Groups   Name        Variance Std.Dev. Corr
```

```
 part_id  (Intercept) 0.62209  0.7887
         present1    4.67694  2.1626   -0.02
 action   (Intercept) 0.41772  0.6463
         present1    0.24248  0.4924    0.96
         expName1    0.07418  0.2724   -0.55 -0.62
Number of obs: 10504, groups:  part_id, 45; action, 19

Fixed effects:
            Estimate Std. Error z value Pr(>|z|)
(Intercept)   1.2478     0.2020   6.179 6.46e-10 ***
expName1      0.4097     0.1433   2.859  0.00425 **
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Correlation of Fixed Effects:
         (Intr)
expName1 -0.047
```

Moral Wrongness:

```
Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
 Family: binomial  ( logit )
Formula: match ~ (1 + present | part_id) + (1 + expName | action) + expName
   Data: subset(tmpdata.1.4, speedCond == "speeded" & judgmentType ==      "moral.Agent")

     AIC      BIC   logLik deviance df.resid
  9176.8   9234.4  -4580.4   9160.8     9810

Scaled residuals:
     Min       1Q   Median       3Q      Max
-10.9451  -0.5280   0.1874   0.5129   4.0717

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 part_id  (Intercept) 0.65064  0.8066
         present1    3.14774  1.7742   0.13
 action   (Intercept) 0.64813  0.8051
         expName1    0.09602  0.3099   -0.06
Number of obs: 9818, groups:  part_id, 86; action, 19

Fixed effects:
            Estimate Std. Error z value Pr(>|z|)
(Intercept)   1.0476     0.2130   4.918 8.74e-07 ***
expName1      0.3910     0.1208   3.236  0.00121 **
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Correlation of Fixed Effects:
         (Intr)
```

```
expName1 0.004
```

Results for Comparisons of Study 4 and Study 3

Color:

```
Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
 Family: binomial  ( logit )
Formula: match ~ (1 + present | part_id) + (1 + expName | action) + present +      expName + present:expName
   Data: subset(tmpdata.3.4, speedCond == "speeded" & judgmentType ==      "color")

     AIC      BIC   logLik deviance df.resid
  7626.5   7699.2  -3803.2   7606.5    10642

Scaled residuals:
    Min      1Q  Median      3Q     Max
-8.7109 -0.0723  0.2040  0.3942  4.4901

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 part_id  (Intercept) 0.53982  0.7347
          present1    2.45963  1.5683   -0.07
 action   (Intercept) 0.06008  0.2451
          expName1    0.06284  0.2507   0.79
Number of obs: 10652, groups:  part_id, 50; action, 19

Fixed effects:
                  Estimate Std. Error z value Pr(>|z|)
(Intercept)         1.6468     0.1341  12.282  < 2e-16 ***
present1            1.1489     0.2332   4.927 8.35e-07 ***
expName1            1.2405     0.1335   9.289  < 2e-16 ***
present1:expName1  -0.8584     0.2325  -3.691 0.000223 ***
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Correlation of Fixed Effects:
            (Intr) prsnt1 expNm1
present1    -0.019
expName1     0.020  0.007
prsnt1:xpN1  0.016 -0.168 -0.044
```

Role:

First, model comparisons revealing no effect of Experiment Name:

```
> anova(lm.0, lm.1h2, lm.2, lm.3)
Data: subset(tmpdata.3.4, speedCond == "speeded" & judgmentType ==  ...
Models:
lm.0: match ~ 1 + (1 + present + harmLevel | part_id) + (1 + present + expName | action)
lm.1h2: match ~ (1 + present + harmLevel | part_id) + (1 + present + expName | action) + harmLevel
lm.2: match ~ (1 + present + harmLevel | part_id) + (1 + present + expName | action) + harmLevel + expName
lm.3: match ~ (1 + present + harmLevel | part_id) + (1 + present + expName | action) + harmLevel + expName
+ harmLevel:expName
       npar    AIC   BIC  logLik deviance  Chisq Df Pr(>Chisq)
lm.0     13 9942.2 10037 -4958.1   9916.2
lm.1h2   14 9938.0 10040 -4955.0   9910.0 6.1711  1    0.01299 *
lm.2     15 9940.0 10050 -4955.0   9910.0 0.0000  1    1.00000
lm.3     16 9941.3 10058 -4954.6   9909.3 0.7325  1    0.39208
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

> anova(lm.0, lm.1h2, lm.3)
Data: subset(tmpdata.3.4, speedCond == "speeded" & judgmentType ==  ...
Models:
lm.0: match ~ 1 + (1 + present + harmLevel | part_id) + (1 + present + expName | action)
lm.1h2: match ~ (1 + present + harmLevel | part_id) + (1 + present + expName | action) + harmLevel
lm.3: match ~ (1 + present + harmLevel | part_id) + (1 + present + expName | action) + harmLevel + expName
+ harmLevel:expName
       npar    AIC   BIC  logLik deviance  Chisq Df Pr(>Chisq)
lm.0     13 9942.2 10037 -4958.1   9916.2
lm.1h2   14 9938.0 10040 -4955.0   9910.0 6.1711  1    0.01299 *
lm.3     16 9941.3 10058 -4954.6   9909.3 0.7325  2    0.69334
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Now, the best-fitting model:

```
Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
 Family: binomial  ( logit )
Formula: match ~ (1 + present + harmLevel | part_id) + (1 + present + expName |      action) + harmLevel
   Data: subset(tmpdata.3.4, speedCond == "speeded" & judgmentType ==      "role")

     AIC      BIC   logLik deviance df.resid
  9938.0  10040.3  -4955.0   9910.0    10976

Scaled residuals:
    Min      1Q  Median      3Q     Max
-6.7568 -0.2719  0.3187  0.4936  4.3935

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 part_id  (Intercept) 0.94458  0.9719
          present1    0.87081  0.9332   -0.15
          harmLevel1  0.09564  0.3093   -0.16 -0.11
 action   (Intercept) 0.22318  0.4724
          present1    0.21722  0.4661   -0.47
```

```
        expName1    0.02154  0.1468    0.09 -0.73
Number of obs: 10990, groups:  part_id, 50; action, 19

Fixed effects:
            Estimate Std. Error z value Pr(>|z|)
(Intercept)   1.5223     0.1749   8.702   <2e-16 ***
harmLevel1   -0.3052     0.1163  -2.624   0.0087 **
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Correlation of Fixed Effects:
          (Intr)
harmLevel1 -0.230
```

Harm:

First, model comparisons revealing a marginal effect of Experiment Name:

```
> anova(lm.0, lm.2e)
Data: subset(tmpdata.3.4, speedCond == "speeded" & judgmentType ==  ...
Models:
lm.0: match ~ 1 + (1 + present | part_id) + (1 + present + expName | action)
lm.2e: match ~ (1 + present | part_id) + (1 + present + expName | action) + expName
      npar   AIC   BIC  logLik deviance  Chisq Df Pr(>Chisq)
lm.0    10 10648 10722 -5314.2    10628
lm.2e   11 10647 10728 -5312.5    10625 3.4276  1    0.06412 .
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Now, the marginally better-fitting model with the variable Experiment Name included:

```
Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
 Family: binomial  ( logit )
Formula: match ~ (1 + present | part_id) + (1 + present + expName | action) +      expName
   Data: subset(tmpdata.3.4, speedCond == "speeded" & judgmentType ==      "harm")

     AIC      BIC   logLik deviance df.resid
 10646.9  10728.2  -5312.5  10624.9    11899

Scaled residuals:
    Min      1Q  Median      3Q     Max
-8.0543 -0.4748  0.2180  0.5211  3.9746

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 part_id  (Intercept) 0.4264   0.6530
          present1    4.0654   2.0163   -0.22
```

```
 action   (Intercept) 0.2856   0.5344
        present1    0.4371   0.6611    0.81
        expName1    0.1263   0.3554   -0.02 -0.20
Number of obs: 11910, groups:  part_id, 54; action, 19

Fixed effects:
          Estimate Std. Error z value Pr(>|z|)
(Intercept)   1.1040     0.2029   5.440 5.32e-08 ***
expName1      0.2401     0.1268   1.894   0.0583 .
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Correlation of Fixed Effects:
        (Intr)
expName1 0.020
```

Moral Wrongness:

First, model comparisons with and without the Experiment Name variable:

```
Data: subset(tmpdata.3.4, speedCond == "speeded" & judgmentType ==  ...
Models:
lm.0: match ~ 1 + (1 + present | part_id) + (1 + expName | action)
lm.2e: match ~ (1 + present | part_id) + (1 + expName | action) + expName
     npar    AIC    BIC  logLik deviance  Chisq Df Pr(>Chisq)
lm.0     7 8895.6 8945.6 -4440.8   8881.6
lm.2e    8 8896.2 8953.4 -4440.1   8880.2 1.3434  1     0.2464
```

Now, the model with the variable Experiment Name included (despite lack of significance):

```
Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']
 Family: binomial  ( logit )
Formula: match ~ (1 + present | part_id) + (1 + expName | action) + expName
   Data: subset(tmpdata.3.4, speedCond == "speeded" & judgmentType ==      "moral.Agent")

     AIC      BIC   logLik deviance df.resid
  8896.2   8953.4  -4440.1   8880.2     9300

Scaled residuals:
     Min       1Q   Median       3Q      Max
-10.6860  -0.5566   0.1865   0.5280   4.1847

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 part_id  (Intercept) 0.3838   0.6195
        present1    3.3269   1.8240   0.06
 action   (Intercept) 0.6318   0.7948
```

```
        expName1    0.1660   0.4074   0.06
Number of obs: 9308, groups:  part_id, 79; action, 19

Fixed effects:
           Estimate Std. Error z value Pr(>|z|)
(Intercept)   0.7990     0.2033   3.931 8.45e-05 ***
expName1      0.1468     0.1250   1.175     0.24
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Correlation of Fixed Effects:
         (Intr)
expName1 0.084
```

References

Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language, 59*(4), 390–412. https://doi.org/10.1016/j.jml.2007.12.005

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language, 68*(3), 255–278. https://doi.org/10.1016/j.jml.2012.11.001

Gray, K., & Wegner, D. M. (2009). Moral typecasting: divergent perceptions of moral agents and moral patients. *Journal of personality and social psychology*, *96*(3), 505–520.